

University of Liège

Gembloux Agro-Bio Tech

Life Cycle Assessment of the Data Processing Infrastructure of the Large Hadron Collider beauty Experiment at CERN

Master's thesis by

Roman Dandoy

submitted for the degree of

Master in Environmental Bioengineering

Academic supervisor: Andrea Di Maria

CERN supervisor: Heinrich Schindler

Academic year 2025–2026

No part of this document may be reproduced by any means without the permission of the author and the academic authority of Gembloux Agro-Bio Tech.

The author alone is responsible for the content of this document.



University of Liège

Gembloux Agro-Bio Tech

Life Cycle Assessment of the Data Processing Infrastructure of the Large Hadron Collider beauty Experiment at CERN

Master's thesis by

Roman Dandoy

submitted for the degree of

Master in Environmental Bioengineering

Academic supervisor: Andrea Di Maria

CERN supervisor: Heinrich Schindler

Academic year 2025–2026

Acknowledgements

This work was carried out at CERN within the EP-LCB-TC. I would like to thank everyone in this section, as each and every one of you contributed to making my time here truly unforgettable. Thank you all for your collaboration, for your help throughout this work, and for the warm welcome you gave me.

Ganz besonders möchte ich dir danken, Heinrich. Ohne dich wäre hier nichts dasselbe gewesen – weder die Qualität meiner Arbeit noch die meines Aufenthalts. All das habe ich zu einem großen Teil dir zu verdanken. Du bist mit Abstand der beste Betreuer, den ich mir hätte wünschen können. Danke für alles!

Grazie Francesco per il tuo indispensabile aiuto durante questo lavoro. Grazie per la tua competenza, il tuo buonumore, il tuo senso dell'umorismo e soprattutto per la tua gentilezza e disponibilità. Ho apprezzato davvero tanto tutto questo :).

Grazie Andrea per il tuo sostegno e i tuoi incoraggiamenti! Dal mio tirocinio di master fino alla tesi, mi hai accompagnato in modo impeccabile, lasciandomi allo stesso tempo la libertà di costruire la mia esperienza personale.

Thank you to everyone else who made this adventure even better: Eric, Augusto, Tony, Alessandro, Emiliano, Laurent, Wiktor, and Jan.

Merci à Léa Bitard et à l'équipe resilioDB de m'avoir permis de mener à bien ce travail grâce à votre précieux logiciel.

Merci à l'équipe « MITSI » de l'EPFL (Eric Xavier Ouvrard et Julia Paolini) pour le temps que vous avez consacré à m'aider à développer cette méthodologie.

Merci aussi à Ophélie, et pardon pour toutes les galères administratives que je t'ai fait subir ! Merci pour tout, j'amène les chocolats.

I also feel that it is important to acknowledge the resources, whether freely available or not, that I had access to throughout this work. I would therefore like to thank openLCA and the people who develop and maintain this software.

Merci à mes amis de Gembloux (Kotolique, Neon's, la fac) et aux 3 autres, qui m'ont toujours entouré et qui m'entoureront toujours (vous vous reconnaîtrez).

Gracias a ti, Mire, por todos esos momentos tan especiales que hemos compartido, por tus consejos, por tus clases de baile y, muy pronto, ¡también por las de español! Gracias por estar ahí y apoyarme cuando estaba agotado, y gracias por todo lo que me has aportado. (Y, por cierto, gracias también por hacerme descubrir La sociedad de la nieve. ¡Qué peliculón! La recomiendo muchísimo). Lo mejor está por venir :)

Merci à ma famille pour son soutien indéfectible durant ces six mois. Merci d'avoir toujours été là, et de toujours l'être lorsque j'en ai besoin. Ce travail, c'est grâce à vous. Merci à mes parents, à mes frères, à ma sœur et à ma grand-mère.

Thank you all. In a way, this work is yours too...

Résumé

Ce travail évalue les impacts environnementaux de la chaîne de traitement des données de l'expérience LHCb au CERN. L'analyse couvre les impacts incorporés et opérationnels durant l'année 2025 du centre de données dit online situé sur site et des centres de données dits offline du Worldwide LHC Computing Grid mobilisés par LHCb. Les impacts sont rapportés au traitement de la chaîne de données LHCb pendant une année de prise de données, avec des indicateurs complémentaires fondés sur la luminosité intégrée et les HS23·heures. Les différentes incertitudes sont ensuite propagées par simulation de Monte-Carlo. Aux paramètres centraux, l'impact climatique du système complet est estimé à 3,100 t CO_2eq, dont 2,137 t pour l'offline et 963 t pour l'online. La simulation de Monte-Carlo donne une moyenne de 3,182 t et un intervalle P5–P95 de 2,644 t to 3,729 t. L'électricité domine l'empreinte offline, tandis que la fabrication du matériel occupe une place importante dans l'online. Le stockage sur disques contribue fortement aux impacts incorporés offline. Les résultats montrent que les priorités sont d'améliorer la caractérisation des équipements et de leur consommation dans les centres offline, de prolonger la durée d'usage lorsque le gain d'efficacité ne compense pas l'impact d'un remplacement, et de suivre conjointement les impacts absolus et les transferts entre catégories environnementales.

Abstract

This work assesses the environmental impacts of the data-processing chain of the LHCb experiment at CERN. It covers the online data centre at Point 8 and the share of the Worldwide LHC Computing Grid's offline data centres used by LHCb. Embodied and operational impacts are combined in an attributional cradle-to-grave assessment using the ecoinvent APOS system model and the Environmental Footprint 3.1 method. The reference period is 2025. Impacts are reported for processing the LHCb data chain during one data-taking year, with additional indicators based on integrated luminosity and HS23·hours. Parametric and representativeness uncertainties are propagated by Monte Carlo simulation. At the central parameter values, the climate-change impact of the complete system is estimated at 3,100 t CO_2eq, of which 2,137 t is attributed to the offline branch and 963 t to the online branch. The Monte Carlo simulation gives a mean of 3,182 t and a P5–P95 interval of 2,644 t to 3,729 t. Electricity dominates the offline footprint, whereas hardware manufacture represents a substantial share of the online footprint. Disk storage also contributes strongly to offline embodied impacts. The results identify three priorities: improving the description of equipment and electricity use across offline data centres; extending equipment use when the efficiency gain from replacement does not offset the manufacturing impact; and monitoring absolute impacts together with burden shifts across environmental categories.

List of acronyms

AHU	Air Handling Unit
AI	Artificial Intelligence
APOS	Allocation at the Point of Substitution
CERN	European Organization for Nuclear Research
CI	Carbon Intensity
CPU	Central Processing Unit
DAQ	Data Acquisition
DQS	Data Quality System
EB	Event Builder
ECS	Experiment Control System
EoL	End of Life
EF	Environmental Footprint
GHG	Greenhouse Gas
GSD	Geometric Standard Deviation
GPU	Graphics Processing Unit
GWP100	Global Warming Potential over 100 years
HDD	Hard Disk Drive
HEP	High-Energy Physics
HLT	High-Level Trigger
HS23	HEPScore23
IT	Information Technology
KPI	Key Performance Indicator
LCA	Life Cycle Assessment
LHC	Large Hadron Collider
LHCb	Large Hadron Collider beauty
LS	Long Shutdown
MC	Monte Carlo
NVMe	Non-Volatile Memory Express
PUE	Power Usage Effectiveness
SMT	Simultaneous Multithreading
SSD	Solid-State Drive
TDP	Thermal Design Power
TFC	Timing and Fast Control
WCT	Wall-clock Time
WEEE	Waste Electrical and Electronic Equipment
WLCG	Worldwide LHC Computing Grid
WUE	Water Usage Effectiveness

Contents

Résumé ii

Abstract iii

List of acronyms iv

List of Figures viii

1 Introduction 1
- 1.1 Climate change, a defining societal challenge . . . 1
- 1.2 CERN and its environmental footprint . . . 1
- 1.3 The LHCb experiment . . . 2
  - 1.3.1 Physics programme and detector . . . 2
  - 1.3.2 Luminosity and the LHC operating cycle . . . 3
  - 1.3.3 From collisions to data: the online system . . . 4
  - 1.3.4 Beyond the site: the offline system and the WLCG . . . 6
  - 1.3.5 Sustainability efforts at the LHCb online data centre . . . 6
- 1.4 Life cycle assessment applied to data centres . . . 6
- 1.5 Objectives of this work . . . 7

2 Materials and methods 8
- 2.1 Scope of the study . . . 8
  - 2.1.1 System under study . . . 8
  - 2.1.2 Life-cycle approach and system model . . . 9
  - 2.1.3 System boundaries . . . 9
  - 2.1.4 Reference period . . . 10
  - 2.1.5 General assumptions and known limitations . . . 10
- 2.2 Functional unit and reference flow . . . 11
  - 2.2.1 Methodological principle . . . 11
  - 2.2.2 Application to the LHCb computing system . . . 11
  - 2.2.3 Methodological choices . . . 11
  - 2.2.4 Implementation in the model . . . 12
- 2.3 Methodology . . . 12
  - 2.3.1 General framework . . . 12
  - 2.3.2 Embodied impacts . . . 14
  - 2.3.3 Operational impacts . . . 21
  - 2.3.4 Uncertainty analysis . . . 28

3 Results 33
- 3.1 Online system . . . 33
  - 3.1.1 Embodied impacts . . . 33
  - 3.1.2 Overall online impacts . . . 35
- 3.2 Offline system . . . 37

3.2.1 Embodied impacts . . . 37
3.2.2 Overall offline impacts . . . 37
3.3 Complete system: multi-impact profile . . . 37
3.4 Online and offline climate-change impacts . . . 39
3.5 Uncertainty analysis results . . . 40
3.5.1 Monte Carlo simulation . . . 40
3.5.2 Contributions to the variance . . . 42
3.5.3 Scenario variants . . . 42
3.6 Comparative analysis: online versus offline computing . . . 43
3.6.1 Shares and carbon intensity of computing . . . 43
3.6.2 Impact per unit of recorded luminosity . . . 44

4 Discussion **45**
4.1 Interpretation of the results . . . 45
4.2 Robustness of the results . . . 47
4.3 Limitations . . . 47
4.3.1 Boundary and missing data . . . 47
4.3.2 Modelling assumptions and comparability . . . 48
4.3.3 Temporal interpretation of the luminosity indicator . . . 48
4.4 Contributions of this work . . . 48
4.5 Possible improvements . . . 49

5 Conclusions **50**

6 AI in this work **51**

References **52**

Appendices **57**
Methodological supplements . . . 58
Impact categories of the EF 3.1 method . . . 64
Results — embodied impacts of the online system . . . 68
Results — operational impacts of the online system . . . 74
Results — offline system . . . 74
Results — multi-impact profile of the complete system . . . 79
Results — climate-change balance and grid geography . . . 81
Results — uncertainty analysis . . . 87
Results — online/offline comparison and derived indicators . . . 92

List of Figures

1.1 Side view of the LHCb detector and its underground cavern . . . 3
1.2 Instantaneous and integrated luminosity foreseen across the LHCb upgrade programme 4
1.3 LHCb experiment project timeline . . . 4
1.4 Overview of the LHCb experiment infrastructure: the detector in its underground cavern and the online data centre on the surface at the far left . . . 5
1.5 The LHCb online data centre . . . 5

2.1 LHCb real-time data flow, from detector readout to offline storage . . . 8
2.2 Life cycle stages considered . . . 9
2.3 Gigabyte G482-Z51 server [38] . . . 16
2.4 Supermicro SYS-621C-TN12R server [39] . . . 17
2.5 Computing work delivered to LHCb by the offline system in 2025 (CERN in Switzerland) 19
2.6 Schematic of an air handling unit . . . 22
2.7 Evolution of the worldwide average annual data-centre PUE [11] . . . 23
2.8 Daily water consumption of the adiabatic cooling system and daily maximum outside temperature at Point 8 in 2025 . . . 24
2.9 Estimated complete-server power per logical core, calculated from processor TDP, hardware-thread count and a confidential vendor-specific uplift factor . . . 26
2.10 Illustrative server power as a function of CPU utilisation, normalised to peak power P_{max}. The curves follow the empirical model of Fan, Weber, and Barroso [47], $P(u) = P_{\text{idle}} + (P_{\text{max}} - P_{\text{idle}})(2u - u^{1.4})$, for two assumed idle-power fractions. The figure illustrates non-proportional behaviour and is not used to derive P_{core} . . . 27

3.1 Annual embodied impacts of the online system by equipment category — EF v3.1 . . . 33
3.2 Embodied climate-change impact of one unit from each server class, by physical subassembly . . . 34
3.3 Embodied climate-change impact of one unit from each server class, by manufacturing process . . . 35
3.4 Total annual impacts of the online system — EF v3.1 . . . 36
3.5 Total annual impacts of the complete system at the central parameter values — EF v3.1 38
3.6 Person-equivalent normalisation of the complete-system EF 3.1 results for 2025, using the EF 3.1 normalisation factors . . . 38
3.7 Deterministic climate-change balance at the central parameter values: (a) Electricity Maps; (b) ecoinvent 3.12 . . . 39
3.8 Operational climate-change impact of LHCb offline computing in 2025 by host country — Electricity Maps . . . 40
3.9 Distribution of the annual climate-change impact of the complete system in 2025 — Electricity Maps annual averages . . . 41
3.10 Monte Carlo climate-change balance by electricity-data source: means by contribution and P5–P95 intervals of the online and offline totals . . . 41
3.11 First-order contributions of the four uncertainty groups to the variance of the climate-change result . . . 42

3.12 Effect of modelling variants on the Monte Carlo mean of the climate-change result in 2025 . . . 43
3.13 Climate-change impact per unit of provisioned CPU capacity — Electricity Maps . . . 44

4.1 Effect of individual drive capacity at constant total provisioned capacity: (a) embodied impact per terabyte and year; (b) electrical power per terabyte of WD Gold drives under operating and idle conditions . . . 46

1 Daily average LHCb electrical power in 2025 . . . 58
2 CERN cable-numbering convention . . . 60
3 Annual embodied climate-change impact of the online system by equipment family . . . 71
4 Online IT climate-change impact by life-cycle stage . . . 71
5 Online server processors by manufacturing process . . . 72
6 Online IT impacts by life-cycle stage — EF 3.1 . . . 73
7 Annual embodied impacts of the offline system — EF 3.1 . . . 75
8 Annual impacts of the offline system — EF 3.1 . . . 77
9 Monthly LHCb grid storage in 2025 . . . 78
10 Freshwater-ecotoxicity impact of a 16 TB grid HDD — ResilioDB . . . 78
11 Complete-system impact profile — EF 3.1, ecoinvent 3.12 APOS . . . 81
12 Online climate-change impacts by component — Electricity Maps . . . 82
13 Offline climate-change impacts by component — Electricity Maps . . . 82
14 Offline computing and climate-impact shares — Electricity Maps . . . 84
15 Offline computing climate-change impact by host country — Electricity Maps . . . 86
16 Grid-computing carbon intensity by host country . . . 86
17 Online and offline Monte Carlo distributions — Electricity Maps . . . 87
18 Leading inventory contributors to climate-result variance . . . 88
19 Complete-system Monte Carlo distributions — EF 3.1, ecoinvent 3.12 APOS . . . 91
20 Online/offline electricity and climate-impact shares — Electricity Maps . . . 92
21 Monthly Point 8 electricity use in 2025 . . . 93
22 Monthly Point 8 PUE in 2025 . . . 94
23 WD Gold delivered price per terabyte in Switzerland on 4 August 2026 . . . 94

1. Introduction

1.1 Climate change, a defining societal challenge

Human societies do not stand outside the environment as we sometimes tend to feel; they are built upon it. The food we eat, the water we drink, the fibres we wear and the materials our infrastructures are made directly or through industrial transformation, from natural systems. Beyond this material perspective, ecosystems provide services that no technology replaces at scale: pollination of crops, purification of freshwater, regulation of the climate, and the cycling of the nutrients on which agriculture depends [1]. Protecting the environment is not an act of altruism towards nature; it is a matter of preserving the conditions of our own prosperity, and ultimately of our own subsistence.

The planetary boundaries framework formalises this dependence by defining, for nine Earth-system processes, the quantified limits within which humanity can operate safely. Recent assessments estimate that six of these nine boundaries are already transgressed, including those for climate change, biosphere integrity and novel entities [2]. Among them, climate change occupies a singular position: it is both one of the most advanced transgressions and a driver that aggravates most of the others. The Intergovernmental Panel on Climate Change (IPCC) states that human influence has unequivocally warmed the atmosphere, oceans and land, and that limiting warming to the levels agreed in Paris requires rapid and sustained reductions of greenhouse gas (GHG) emissions across every sector of activity [3].

Science is not exempt from this global reduction effort. Its increasing reliance on digital technologies makes computing a growing part of its environmental footprint: data centres alone consumed an estimated 415 TW h of electricity in 2024 — around 1.5 % of global demand — and this consumption is expected to more than double by 2030 under the combined pressure of cloud services and artificial intelligence [4].

The impacts of computing are moreover not limited to the electricity it draws: manufacturing servers is a resource- and energy-intensive industrial process, whose burden is incurred before the first job ever runs. The scientific community has begun to examine its own footprint — and high-energy physics (HEP), which operates some of the largest research infrastructures and computing systems in the world, is no exception [5]. This work is a contribution to that effort, applied to the data centres of one particular experiment: LHCb, at CERN.

1.2 CERN and its environmental footprint

The European Organization for Nuclear Research (CERN), founded in 1954 and located on the Franco-Swiss border near Geneva, is the largest particle physics laboratory in the world. Its main mission is fundamental research: probing the elementary constituents of matter and the forces that govern them. To this end, CERN operates a chain of particle accelerators culminating in the Large Hadron Collider (LHC), a 27 km ring buried about 100 m underground, in which protons (and occasionally heavy ions) are collided at energies of up to 13.6 TeV. Four large experiments are installed at the collision points of the ring — ATLAS, CMS, ALICE and LHCb — each operated by an international collaboration of thousands of scientists.

Since 2020, CERN has published biennial environment reports documenting its energy use, water consumption, direct emissions and mitigation objectives [6]. Although CERN's electricity is supplied by the French grid — whose largely nuclear and hydroelectric mix gives it one of the lowest carbon intensities in Europe — electricity consumption nevertheless remains, according to the environment report, the main contributor to CERN's carbon emissions [6]. In 2025, CERN consumed 1.285 TW h for all of its on-site activities, split between the Super Proton Synchrotron (SPS) and the Prévessin site (31.52 %), the Proton Synchrotron (PS) and the Meyrin site (14.68 %), and the LHC (53.79 %) [7]. Still in 2025, amongst the total energy consumption of LHC, only 21.14 % has been dedicated to the operation of the four large LHC experiments.

Computing occupies a specific position, accelerators and detector magnets dominate instantaneous power draw, but they only draw it when beams circulate, whereas the data centres that process the collisions run essentially year-round. It reprensent an almost permanent operational load and a regularly renewed hardware — that motivates a dedicated environmental assessment of the computing infrastructure, conducted here for the LHCb experiment.

1.3 The LHCb experiment

1.3.1 Physics programme and detector

LHCb (Large Hadron Collider beauty) is one of the four large LHC experiments. It is dedicated to the physics of the beauty quark: by measuring with high precision the decays of hadrons containing b and c quarks, the collaboration studies the subtle differences between matter and antimatter, the violation of the charge-parity (CP) symmetry and searches for indirect manifestations of physics beyond the Standard Model. One of the driving questions of this programme is among the most fundamental in physics: why the observable universe is made of matter, when the Big Bang should have produced matter and antimatter in equal amounts.

Because $b\bar{b}$ pairs are produced predominantly at small angles with respect to the beam axis, LHCb is built as a single-arm forward spectrometer rather than as a cylindrical detector enclosing the collision point (such as CMS or ATLAS detectors). Installed in a cavern approximately 100 m underground at Point 8 of the LHC ring, the detector is a sequence of specialised subsystems (figure 1.1): a vertex locator surrounding the interaction region, ring-imaging Cherenkov detectors for particle identification, tracking stations and a dipole magnet to measure particle momenta, calorimeters to measure energies, and muon chambers at the far end. Each proton-bunch crossing produces signals in these subsystems; reconstructing from those signals the trajectories, identities and decay vertices of the particles is the computational task around which this work was carried out.

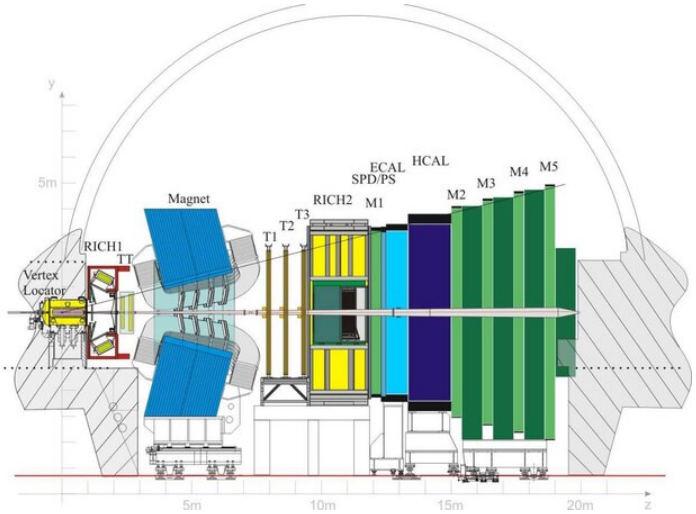


Figure 1.1: Side view of the LHCb detector and its underground cavern

The next major upgrade (Upgrade II) is planned for Long Shutdown 4, in the mid-2030s, to exploit the High-Luminosity LHC [8].

1.3.2 Luminosity and the LHC operating cycle

Two quantities of accelerator physics are needed to understand the computing load of the experiment — and hence the environmental questions addressed later in this work.

The *instantaneous luminosity* $\mathcal{L}$, expressed in $1/(\mathrm{cm}^2\,\mathrm{s})$, characterises the collision rate a collider delivers. For two counter-rotating beams of N_1 and N_2 protons per bunch, colliding at a frequency f with transverse beam sizes σ_x and σ_y, it reads

$$\mathcal{L} = f\,\frac{N_1 N_2}{4\pi\,\sigma_x \sigma_y}, \tag{1.1}$$

so that the number of interactions of a given process per unit time is $\mathcal{L}\sigma$, where σ is the cross section of that process. The *integrated luminosity*, expressed in inverse femtobarns [fb^{-1}], is the accumulation of $\mathcal{L}$ over time and is directly proportional to the total number of collisions recorded. Integrated luminosity is therefore the quantity in which an LHC experiment expresses the size of the data sample underpinning its physics results, which is why it will reappear in this work as a candidate for normalising environmental impacts.

During Run 3, LHCb operates at an instantaneous luminosity of about 2×10^{33} $1/(\mathrm{cm}^2\,\mathrm{s})$. The limit here is not the accelerator, which delivers far more: it is the rate the detector and its computing can absorb without degrading the reconstruction. In the Low scenario retained here, Upgrade II raises that ceiling to 1×10^{34} $1/(\mathrm{cm}^2\,\mathrm{s})$, about five times the Run 3 operating point (figure 1.2) [8]. This increase translates into more simultaneous interactions per bunch crossing, more complex events, and therefore substantially larger data flows to process — with direct consequences for the size and power draw of the computing infrastructure studied here.

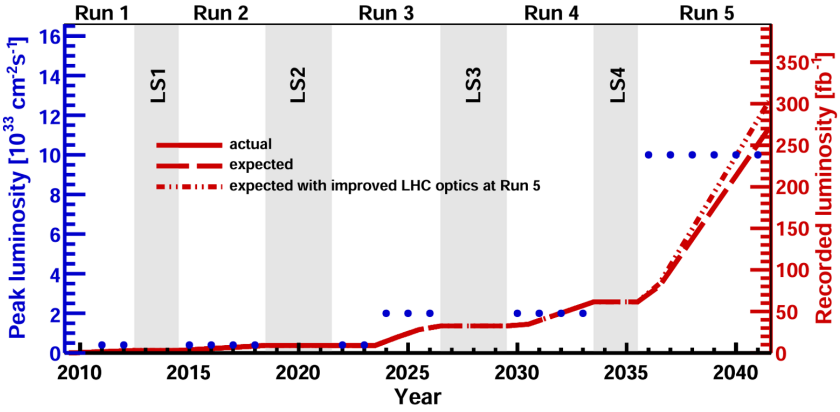


Figure 1.2: Instantaneous and integrated luminosity foreseen across the LHCb upgrade programme

The LHC does not collide protons continuously. Data-taking periods (*Runs*), lasting several years, alternate with *Long Shutdowns* (LS) of one to three years devoted to maintenance and upgrades of the accelerator complex and detectors, with shorter technical stops in between; figure 1.3 situates the successive periods of the experiment, which is at present in LS3.

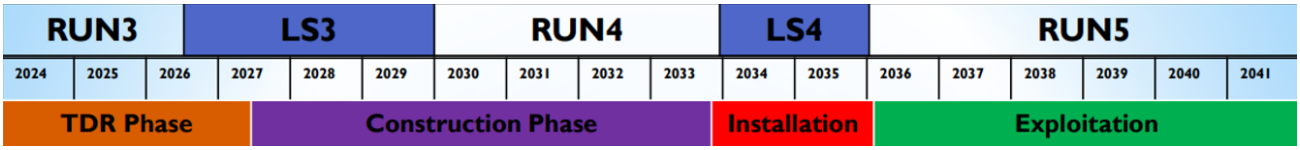


Figure 1.3: LHCb experiment project timeline

This cycle matters for the present work: during a RUN year the LHCb online data centre processes live collision data, whereas during an LS year no collisions occur — yet the servers do not stand idle. They are made available to the collaboration for other workloads, chiefly Monte Carlo production for offline physics analyses. A RUN year and an LS year therefore constitute two operating regimes of the same infrastructure, as discussed in the functional-unit section (section 2.2).

1.3.3 From collisions to data: the online system

Since Upgrade I, LHCb operates without any hardware trigger. Bunches cross in the detector at the 40 MHz clock rate of the LHC, but only some of those will deposit signals worth reading; at the Run 3 operating point this amounts to some 30 MHz. Every one of them is read out in full and processed in software. The detector sends some 4 TB of data per second — an internet-scale traffic — over optical links to roughly 600 custom PCIe40 (TELL40) readout boards. These boards are not standalone units: they are PCIe cards hosted inside the servers of the *event builder*, the farm of interconnected machines whose task is to gather, for each collision, the fragments produced by every subdetector and to assemble them into a single complete *event* — the object on which all subsequent processing operates [9].

This enormous amount of data is then reduced in two stages by the High Level Trigger (HLT). The first stage, HLT1, performs a fast partial reconstruction of every event and discards about 97 % of them; this massively parallel workload runs on graphics processing units (GPUs). Surviving events are written to a disk buffer while the detector is aligned and calibrated in near-real time; the second stage, HLT2, then performs the full reconstruction on a large farm of CPU servers and selects the events to be kept, delivering of the order of 10 GB/s to permanent storage [9]. Figure 2.1 summarises this chain. Together, the event builder and the two trigger stages constitute

what the collaboration calls the *online* system and every single step explained so far takes place on site, in France, at the LHCb experiment.

Physically, the online system does not sit underground with the detector. The data are brought to the surface by optical fibres and processed in a modular container data centre located directly above, at Point 8 (figures 1.4 and 1.5).

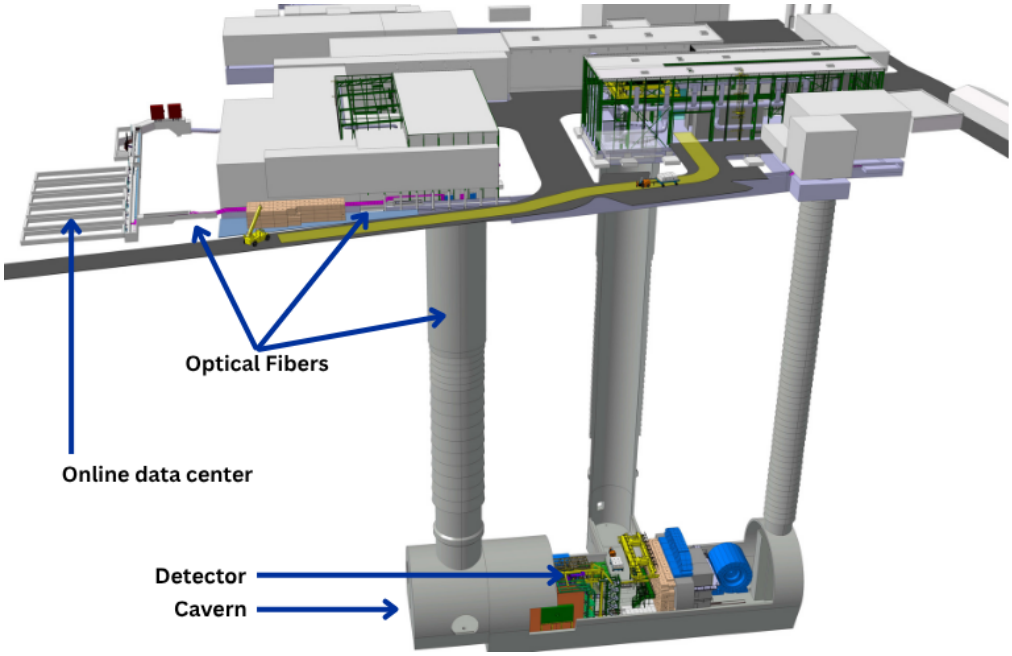


Figure 1.4: Overview of the LHCb experiment infrastructure: the detector in its underground cavern and the online data centre on the surface at the far left

The facility comprises six *IT modules* and one *power module*. All the computing operations are cooled by indirect free-air cooling with adiabatic assistance: outside air removes the heat through the exchangers, and when it becomes too warm to do so on its own, water is sprayed onto the incoming air stream to lower its temperature. Altogether, the modules house approximately 4,600 servers, whose composition is detailed in section 2.3.2.



Figure 1.5: The LHCb online data centre

In 2025, the online data centre accounted for 25.99 % of LHCb's electricity consumption; its metered daily power profile is shown in figure 1, in the appendix.

1.3.4 Beyond the site: the offline system and the WLCG

The 10 GB/s data leaving HLT2 is distributed over the Worldwide LHC Computing Grid (WLCG) for centralised reprocessing campaigns, Monte Carlo simulation, user analyses and long-term storage on disk/tape. The WLCG is a federation of computing centres spanning dozens of countries whose members commit resources to the experiments through pledges [10]. The grid is organised in hierarchical tiers: the Tier-0 at CERN, national Tier-1 centres providing storage and reprocessing capacity, and Tier-2 centres primarily devoted to simulation and analysis. For LHCb, this *offline* system is composed of some forty Tier-0, Tier-1 and Tier-2 sites in thirteen countries. This second part of the data flow chain is also considered in this work.

1.3.5 Sustainability efforts at the LHCb online data centre

The LHCb online data centre already incorporates several efficiency measures. Its free-air cooling architecture yields a Power Usage Effectiveness (PUE) — the ratio of total facility energy to the energy delivered to IT equipment — of 1.05 over 2025, compared with a worldwide mean of approximately 1.54 [11]. Within the ODISSEE project, an artificial-intelligence system is being deployed to optimise cooling operations, notably AHU fan speeds, in order to lower the absolute energy use of the cooling infrastructure [12]. Direct liquid cooling is also under evaluation [13].

1.4 Life cycle assessment applied to data centres

Life Cycle Assessment (LCA) is a standardised methodology for quantifying the environmental impacts of a product, service or system over its entire life cycle — from raw material extraction through manufacturing, transport and use, to end of life. Its framework is codified by the ISO 14040 and 14044 standards, which structure any study into four iterative phases: definition of the goal and scope, life cycle inventory, impact assessment, and interpretation [14, 15]. The multi-criteria approach is relevant here because it quantifies the carbon footprint but also the impact on water use, mineral and metal resources, ecotoxicity and a dozen other categories, thereby guarding against solutions that merely displace a burden from one impact to another.

The data centre industry does track environmental performance, but only through key performance indicators (KPIs) such as the PUE or the Water Usage Effectiveness (WUE), standardised in the ISO/IEC 30134 series [16, 17]. These indicators are valuable operational dashboards, yet they are not impact measures: a facility can display an excellent PUE while running carbon-intensive electricity through recently manufactured, rapidly renewed hardware. Their status with respect to the LCA framework is discussed in section 2.2. Sector-specific methodological guidance does exist, notably the ITU-T L.1410 recommendation for ICT goods and services [18], but its application to a scientific computing infrastructure remains far from routine.

The main challenge addressed in this thesis is the availability and quality of data. Manufacturers of servers and electronic components rarely disclose bills of materials or environmental product declarations. Consequently, both the composition of a specific server model and the environmental impacts associated with its production must be modelled from generic databases, whose electronics datasets are often highly aggregated, too generic, and often outdated. Several research communities are working to address these limitations through initiatives such as parametric electronics databases, including ResilioDB; bottom-up models of integrated-circuit manufacturing [19]; environmental assessments of computing infrastructures at the scale of scientific experiments, such as the recent ATLAS study [20]; and measurement campaigns of power consumption and energy

efficiency at computing sites within CERN and the HEPiX community [21, 22]. Yet estimating the embodied footprint of computing hardware remains a major challenge, and much of the available literature on sustainable scientific computing still takes the form of recommendations and good practices rather than quantified and comprehensive assessments of systems [23, 24].

One further constraint distinguishes scientific data centres from their commercial counterparts and narrows the space of applicable solutions: the performance they must deliver is not negotiable. The online farm must absorb 4 TB/s in real time — every second of beam lost to a computing failure is physics lost — and the grid must keep pace with the experiment's data production. Sustainability levers must therefore be sought within, not against, these throughput requirements; this tension between environmental ambition and operational necessity runs through the entire work.

1.5 Objectives of this work

The objective of this work is to identify, as reliably as possible, the main contributors to the environmental impacts of the LHCb experiment's data centres. It aims to quantify their relevant environmental impacts — embodied and operational, online and offline — and to identify improvement measures for the subsystems analysed.

More specifically, this work sets out to answer the following questions:

1. What are the embodied and operational environmental impacts of the computing chain of the LHCb experiment — the online data centre and WLCG resources — over one reference year, under RUN operating conditions?
2. Which subsystems, components and life cycle stages dominate each impact category, and where do the greatest reduction opportunities therefore lie?
3. How do the online and offline parts of the chain compare, given that the first is assessed from an exhaustive hardware inventory and the second from top-down estimates?
4. How robust are these results to the uncertainties (data, models and methodological scenarios) regarding the assessment of computing hardware?
5. Which improvement levers are realistic given the performance requirements the infrastructure must meet?

This document is intended first and foremost to provide the LHCb collaboration with a quantified reference on the environmental impacts of its computing activities. It also seeks, at its own modest scale, to bring together three communities: sustainability science (through the LCA method), information and communication technology (the data centres) and high-energy physics (LHCb). The concepts specific to each field are thus presented as the work progresses. Beyond the numbers this work provides for a specific case, it aims to leave behind a methodology that others can analyse and, potentially, apply to another complex system.

2. Materials and methods

2.1 Scope of the study

2.1.1 System under study

The system under study is the computing infrastructure that processes data from the LHCb experiment. This infrastructure is divided into two parts.

The *online* part, already briefly described in the introduction, is the data centre located on site, at the surface of Point 8. This physical facility consists of six IT modules and one power module, an air-handling system, an electrical distribution chain and a fleet of computing servers, together with the racks, network equipment and cables that connect them and the optical fibres that bring detector data up from the cavern.

The *offline* part is the share of the Worldwide LHC Computing Grid used by LHCb: the Tier-0, Tier-1 and Tier-2 sites that execute the experiment's reprocessing, simulation and analysis workloads, and that hold its data on disk and on tape. This infrastructure does not belong to the collaboration, and only part of its hardware is inventoried.

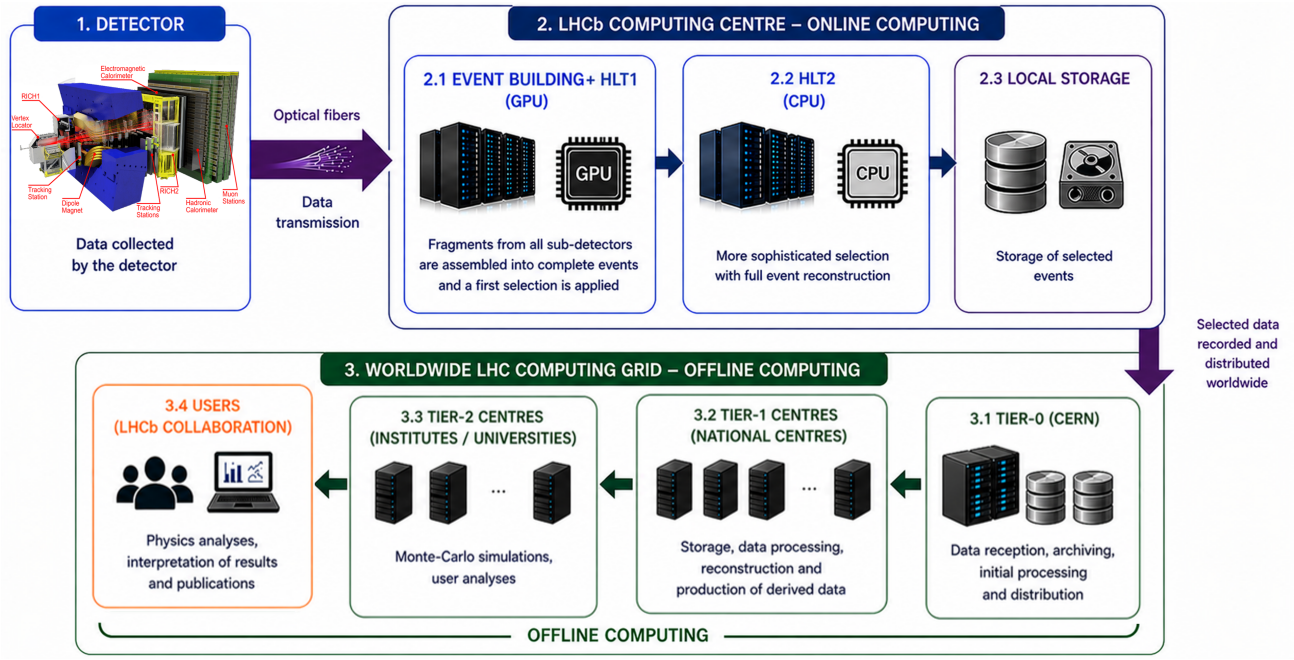


Figure 2.1: LHCb real-time data flow, from detector readout to offline storage

This difference in the depth of coverage is not a gap that a little more data collection would close. Closing it would mean instrumenting and reporting the hardware of some forty independent sites, each with its own operator, procurement cycle and monitoring practice. Such an undertaking is beyond the scope of the present work, though it would improve the precision of any assessment of the grid. This gap was a major challenge throughout this work.

What lies outside the system is the experiment itself. The detector, its front-end electronics, the cavern, the accelerator and the electricity they consume are not assessed: this work concerns the processing of the data, not their production. The boundary is drawn at the PCIe40 readout boards,

which are counted as part of the online system since they sit in the event builder servers.

2.1.2 Life-cycle approach and system model

The assessment is conducted from a *cradle-to-grave* perspective: for every element retained in the boundary, the impacts are counted from the extraction of the raw materials through manufacturing, transport, installation and use, to end-of-life treatment.

The foreground inventory was implemented in openLCA 2.6.2 and ResilioDB 2026.7. Background processes are drawn from ecoinvent 3.12 using its APOS system model. Operational carbon intensities use the Electricity Maps 2025 yearly dataset as redistributed in version 0.18.0 of CO2.js, the Green Web Foundation's `@tgwf/co2` software package; recording these versions makes the calculations reproducible even when the databases are updated [25–29].

The assessment follows an attributional approach. Among the attributional system models available in ecoinvent, the allocation at the point of substitution (APOS) model was selected for the background system. This choice was made to remain consistent with the foreground modelling, where waste treatment and end-of-life processes are explicitly assigned to the system generating the waste. Using APOS for the background therefore ensures a consistent treatment of waste flows throughout the assessment.

A consolidated overview of the foreground data, background models and outputs of the four assessment blocks is provided in table 1.

2.1.3 System boundaries

The life cycle stages retained are those shown in figure 2.2. *Data centre construction* and *data centre deconstruction* mean the assembly and disassembly of the equipment after production and transport to the site. Transport from the data centre to waste-treatment facilities is not modelled.

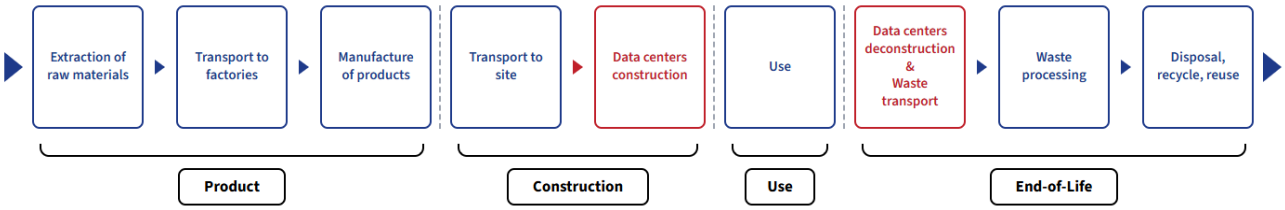


Figure 2.2: Life cycle stages considered

For the physical elements retained within each stage, the two branches (online, offline) do not have the same coverage, for the reason given above. The online system is assessed exhaustively: servers, network equipment, cables, racks, the air handling units, the electrical infrastructure, the containers and their foundations. The offline system is assessed on its compute servers and its disk storage only; magnetic tape, the infrastructures of the grid sites, and the inter-site transfer networks fall outside the boundary Table 2.1.

Table 2.1: Elements excluded from the system boundary

Excluded element	Branch	Grounds
Detector, cavern and accelerator	Both	Outside the function studied: this work assesses the processing of the data, not their production
Staff commuting and business travel	Both	No robust allocation factor (section 2.3.3)
Data centre lighting	Both	Contribution judged negligible
Software development and end-user devices	Both	Shared across the whole collaboration and considered negligible
Tier-3 sites	Offline	Bound by no pledge, and not reported in the centralised accounting; marginal share of the computing volume (section 2.3.1)
Grid site buildings and technical installations	Offline, embodied	Not modellable across some forty heterogeneous sites; their electricity is retained through the PUE (sections 2.3.2 and 2.3.3)
Tape storage	Offline	No literature allowing a reliable impact to be built (section 2.3.2)
Water consumption of the grid sites	Offline, operational	Heterogeneous cooling technologies, and meters covering whole shared facilities (section 2.3.3)
Inter-site transfer networks	Offline	Shared among many users; the negligibility of the LHCb share is discussed (section 2.3.2)

2.1.4 Reference period

The assessment is carried out for a single reference year, 2025. It is the most recent complete year available. The on-site metering, the grid accounting and the carbon intensity series all cover it in full. A single year is also a convenient period over which to bring the two terms together: the embodied impacts, divided by the lifespan assigned to each element, and the operational impacts of the year itself, following equations (2.1) and (2.2). Both then refer to the same period. Finally, in order to make relevant statements, it makes more sense to use the configuration which is the most up-to-date.

2.1.5 General assumptions and known limitations

Under the steady-state assumption, the offline workload performed in 2025 includes the processing of data acquired in previous years, while some of the data acquired the same year will only be processed in subsequent years. The analysis therefore attributes computational work to the year in which it is performed, rather than to the year in which the corresponding data are acquired. This annual accounting approach provides an accurate representation of the workload associated with a typical year if computing-grid activity remains approximately stable over time.

A second limitation concerns data on electronics and is documentary rather than methodological. Server manufacturers publish neither bills of materials nor, save for a few exceptions, environmental product declarations, so the manufacturing impact of a given machine has to be reconstructed from generic or parametric databases whose electronics datasets are aggregated and, in places, dated.

2.2 Functional unit and reference flow

2.2.1 Methodological principle

The purpose of an LCA is to quantify one or more environmental impacts induced by a process, a service, or a product. Its aim is to support the management of the environmental impacts associated with human activities by assessing them in detail. Therefore, in order to choose one system, process, or product over another, it is essential to ensure that they are comparable. This is where the functional unit comes in, defined as:

> *"A quantitative description of the function or service for which the assessment is performed, and the basis of determining the reference flow of product that scales the data collection in the next LCA phase, the inventory analysis."* [30]

The functional unit must therefore describe a quantifiable service and provide a meaningful basis for determining the reference flows used in the inventory analysis.

2.2.2 Application to the LHCb computing system

Defining this service is challenging for a scientific computing infrastructure. Common data-centre indicators such as PUE and WUE describe infrastructure efficiency rather than the service delivered [31]. A data centre may, for example, perform Monte Carlo simulations for particle-physics research or process a comparable amount of data for agronomic analyses. Although these activities may involve similar computing resources, they fulfil fundamentally different functions. Neither the amount of data processed nor the energy consumed therefore provides, by itself, a meaningful basis for comparing the services delivered.

HEPScore23 (HS23) is a hardware-independent CPU benchmark used by the high-energy-physics community to express computing capacity. Computing work, expressed in HS23·hours — the product of HS23 capacity and execution time — is hardware-independent but does not measure the scientific usefulness of that work [32]. Integrated luminosity is meaningful within particle physics, but it corresponds neither to a fixed data volume nor to a fixed computing workload. Moreover, because experiments study different physical processes under different operating conditions, one fb^{-1} does not represent the same scientific service across experiments.

2.2.3 Methodological choices

The purpose of this work is to inform the LHCb collaboration about the environmental impacts of its data-processing chain and to identify potential reduction levers, rather than to compare LHCb with other experiments or data centres. The function considered is therefore the computational processing of LHCb data, from online acquisition and real-time event selection to offline processing, simulation, analysis and storage on the WLCG.

The selected functional unit is:

> *Ensuring the computational processing of data from the LHCb experiment over one reference year, under RUN operating conditions.*

Only RUN operating conditions are considered. During Long Shutdowns, the online servers are used for other workloads, principally Monte Carlo production, but the available metering and reporting data for the two branches do not allow results to be reported with the same scope and accuracy for these periods. The implications of a shutdown year are discussed in Discussion.

Both the online and offline systems are included. Although the offline assessment is less detailed and concerns infrastructure largely outside the direct control of the collaboration, excluding it would provide an incomplete representation of the function and could conceal transfers of environmental burdens from the online system to the grid. The two contributions are therefore reported separately, and the greater uncertainty of the offline branch is discussed in section 2.3.4.

2.2.4 Implementation in the model

The reference flows associated with the functional unit comprise the hardware fleet of the Event Filter Farm and the electricity it consumes during the reference year. For the offline branch, they comprise the logical-core time occupied by LHCb on the WLCG during the same period, expressed as Wall-Clock Time by site, together with the disk capacity provisioned for the experiment. These quantities allocate shares of the reference servers and hard drives to LHCb for embodied impacts and are converted into electricity use for operational impacts.

Offline computing work is attributed to the year in which it is performed, irrespective of when the corresponding data were acquired. The model therefore assumes that the workload performed during the reference year is representative of the computing activity associated with a typical year. The implications of this annual accounting approach for the temporal interpretation of the results are discussed in Discussion.

The model reports the environmental impacts of the complete system, covering both the online and offline branches and both embodied and operational impacts, over one reference year. In addition to this main result, two climate-change intensity indicators are calculated. The first, expressed in tonnes of CO_2eq per fb^{-1}, relates the full-chain impact to the integrated luminosity recorded during the reference year. The second, expressed in t CO_2eq/(HS23·h), provides a common denominator for comparing the online and offline computing branches in section 3.6. HS23·hours is used only as a reporting denominator; the embodied impacts of the offline computing equipment are allocated according to the occupation of logical cores, i.e. the job slots exposed to the batch system, as described in section 2.3.2.

2.3 Methodology

2.3.1 General framework

Embodied and operational impacts

The environmental impacts of data centres are differentiated into two classes throughout this work: *embodied* impacts and *operational* impacts.

Operational impacts designate the impacts of the use phase, i.e. the electricity consumed for computing, storage and networking, as well as the associated water consumption. The impacts related to the mobilisation of the human workforce, notably through transport, could arguably be included as well; they are excluded from the system boundary for the reasons developed in section 2.3.3.

Embodied impacts group all the impacts related to the physical hardware: raw-material extraction, component manufacturing, transport, assembly, replacement and end of life.

For a given reference year and impact category j, the total impact is written:

$$I_{\text{tot},j} = I_{\text{emb},j} + I_{\text{op},j} \tag{2.1}$$

where $I_{\text{emb},j}$ is the embodied impact amortised over the reference year and $I_{\text{op},j}$ the impact of that year's operation. The methodologies used to estimate I_{emb} and I_{op} are detailed in sections 2.3.2 and 2.3.3 respectively.

Life cycle impact assessment and interpretation

Environmental interventions in the inventory are characterised with the Environmental Footprint (EF) 3.1 method, selected as the European Commission reference framework for Product and Organisation Environmental Footprint studies and because it reports sixteen environmental impact categories. Midpoint indicators characterise impacts at an intermediate point of the cause–effect chain, whereas endpoint indicators translate them into damages to human health, ecosystems or resources. The midpoint level is retained because it limits the additional modelling and value choices required for endpoint aggregation or a single score. Among the sixteen midpoint categories, climate change is characterised as the Global Warming Potential over a 100-year time horizon (GWP100). Absolute characterised results remain the primary results of this work; person-equivalent normalisation expresses each category relative to the annual impact of an average person, providing a common order-of-magnitude scale across categories, and is used only as an aid to interpretation [33–35].

Both branches are assessed over the complete set of categories j of the EF v3.1 method. The results are, however, produced from two different sources of electricity data. The multi-impact profile rests on the ecoinvent country mixes, the only source that documents the sixteen categories; the climate indicator is additionally computed from the Electricity Maps annual averages for the studied year, which cover that category alone.

The reason for maintaining both is that the ecoinvent mixes mainly represent 2022, whereas electricity generation can change substantially between years. Their composition affects every category, not only climate change: the shares of nuclear, renewable and fossil generation also influence ionising radiation, mineral resources and toxicity. The difference between the ecoinvent and Electricity Maps climate results is therefore reported as a sensitivity scenario (section 2.3.4). Because the two sources also differ in modelling method and system boundary, this comparison measures their combined effect rather than isolating a temporal bias, and it is not extrapolated quantitatively to the other categories.

The WLCG grid: scope and accounting data

The offline part of the LHCb data processing is distributed over the Worldwide LHC Computing Grid (WLCG), whose hierarchical architecture was presented in section 1.3.4. The perimeter retained covers the Tier 0, 1 and 2 sites serving LHCb, i.e. some forty sites spread over thirteen countries. Tier-3 sites are excluded: unlike the first three levels, they are bound to the WLCG by no formal resource commitment (*pledge*) and their activity is consequently not systematically reported in the centralised accounting [10]; their contribution to the total computing volume is in any case marginal compared with that of the pledged sites.

The heterogeneity of the processors deployed on the network makes it impossible to compare sites based solely on their computing hours: one hour on a recent core does not represent the same amount of work as one hour on a ten-year-old core. This is what the HEPScore23 benchmark introduced in section 2.2 normalises.

The activity data used in this work come from the EGI Accounting Portal, accessed through Grafana [36]. Two time metrics are to be understood, and defined in this work as follows:

CPU time [logical cores $\times$ hours]. The cumulated time during which the logical cores allocated to a job actually execute that job's instructions. For a parallel or multi-core job, the times spent on the different cores are added together. Periods during which the job waits for data, for input/output or for a resource are normally not counted as CPU time.

Wall-clock time (WCT) [logical cores $\times$ hours]. The time actually elapsed between the start and the end of a job, as an external clock would measure it, multiplied by the number of logical cores the job reserved.

The portal distinguishes raw wall-clock time from raw wall-clock work. Because the metric used here multiplies elapsed time by the number of reserved logical cores, it corresponds technically to wall-clock work; the WCT label is retained to match the dashboard terminology [36].

The WCT is therefore by definition larger than the CPU time, the gap between the two reflecting the CPU efficiency of the jobs. This efficiency is high for LHCb workloads: the available job-level analysis reports an average CPU utilisation of 96 %, corresponding to a difference of about 4 % between CPU time and WCT [37].

2.3.2 Embodied impacts

Online: modelling approach and databases

The components of the online system were identified, characterised and counted with the help of the online team. Each component modelled was assigned an estimated lifespan within the product system (table 2.2). In order to compare the impacts of those components consistently, their impact must be normalised to the duration of the reference flow — one year. Each component therefore follows:

$$\text{Annual impact} = \frac{\text{Total life cycle impact}}{\text{Lifespan}} \tag{2.2}$$

For components whose lifespan was judged to exceed 50 years and for which a precise value could not be supported, impacts were normalised using a lifespan of 50 years. These lifespan assumptions are addressed in the uncertainty analysis (section 2.3.4).

Table 2.2: Number of items inventoried and lifespans assumed for the components of the online system

Component	Items	Life [yr]	Component	Items	Life [yr]
Air filters	192	0.5	Routers (switches)	318	10
Air handling units	24	20	Servers, class A (Gigabyte)	200	10
Water pumps	24	20	Servers, class B (Supermicro)	200	10
Reverse osmosis pumps	4	20	Servers, class C (generic)	1,225	10
Inverters	144	20	InfiniBand controllers (HLT1)	600	10
Transformers	2	50	HLT1 storage	3,200	10
Racks	132	50	HLT2 storage	720	10
Containers	7	50	Calibration storage	600	10
Foundations (concrete)	—	50	Service storage (NVMe)	48	10
Cables (ECS, signal, TFC)	—	50			

Cables and foundations are not counted in units; their modelling is documented in Cables and Modelling of the remaining online infrastructure.

Two modelling tools were used. Non-IT components were modelled in openLCA 2.6.2 using ecoinvent 3.12. Ecoinvent provides datasets for a broad range of materials and industrial processes and was therefore used as the default database. However, its IT datasets are very generic and outdated, and do not represent the diversity of servers, processors, storage devices and network equipment included in the system. Reconstructing each IT component from generic ecoinvent processes would require detailed information on thousands of individual parts. IT hardware was therefore modelled in ResilioDB 2026.7. ResilioDB provides parameterised models specifically developed for IT equipment and documents the assumptions, limitations and confidence associated with its estimates. Servers, HDDs and routers were consequently modelled using ResilioDB.

ResilioDB also bases its models on ecoinvent, so the two tools share the same characterisation and normalisation factors; the difference lies in the more detailed process modelling.

Only the servers and storage modelled are presented in the remainder of this subsection. The modelling of the other components, including air-handling units, foundations, containers and transformers, is documented in Modelling of the remaining online infrastructure.

Servers

Servers are understood here as the machines performing the *event building* (EB) and running the two software trigger stages (HLT1 and HLT2), excluding network equipment (routers and switches), which is inventoried separately in table 3. All of them are rack servers, i.e. machines mounted in standardised metal cabinets whose usable height is counted in rack units (1U = 44.45 mm).

The total number of installed servers changes frequently with maintenance, decommissioning and procurement: at the time of data collection (2026), the fleet comprised about 1,400 chassis, i.e. nearly 4,400 compute nodes (the distinction is clarified in section 2.3.2).

A node denotes here an autonomous computing unit — with its own motherboard, processors, memory and operating system — sharing with the other nodes of the same chassis only the enclosure, the power supplies and the cooling. From a functional standpoint, a node therefore behaves as an individual server: classes A and B have one node per chassis, whereas class C packs four nodes per high-density 2U chassis.

Modelling the server fleet faces several obstacles. The first is the heterogeneity of the fleet: most of the machines (class C, table 2.3) mix generations, manufacturers and configurations, so that an exhaustive census would be as tedious as the modelling that would follow. Added to this are the regular changes in the number of installed machines and, above all, the lack of knowledge of their usage history: the actual age and condition of each server, and hence its remaining lifetime potential, are not documented. Projecting the exact state of the fleet over the coming years would therefore have required piling up unverifiable assumptions, to the detriment of the representativeness of the results.

Despite the heterogeneity of the fleet, two specific configurations (chassis, processors, memory, storage) are present in significant numbers, enough to define them as classes A and B. A third, generic class groups the rest of the fleet around an approximate configuration of the remaining nodes surveyed on site (table 2.3).

Table 2.3: Server classes (unless otherwise stated quantities are given *per node*)

	Class A (EB/HLT1)	Class B (HLT2)	Class C (HLT2, generic)
Chassis	Gigabyte G482-Z51	Supermicro SYS-621C-TN12R	Multiple
Height (Chassis)	4U	2U	2U
Nodes/Chassis	1	1	4
Chassis	∼180 (+20 spares)	∼200	∼1,225
Nodes (total)	∼180 (+20 spares)	∼200	∼4,901
CPU	2 × AMD EPYC 7502 (32 cores)	2 × Intel Xeon Platinum 8592+ (64 cores)	2 × Intel Xeon E5-2630/2650/2680 v4 (10–14 cores)
GPU	3 × NVIDIA RTX A5000	—	—
FPGA	3 × PCIe40 (TELL40)	—	—
Memory	16 × 32 GB DDR4	16 × 32 GB DDR5	8 × 16 GB DDR4
Storage	1 × 960 GB SATA SSD	2 × 1.92 TB NVMe SSD + 1 × 480 GB NVMe SSD	1 × 960 GB SATA SSD
Network	2 × InfiniBand ConnectX-6	not documented	not documented
Power supplies (/chassis)	3 × 2,200 W	2 × 1,200 W	2 × 1,600 W (shared)
Fleet total	∼1,625 chassis, ∼5,301 nodes		



(a) Front view (b) Top view

Figure 2.3: Gigabyte G482-Z51 server [38]



(a) Front view (b) Top view

Figure 2.4: Supermicro SYS-621C-TN12R server [39]

The GPUs in the class A servers are better suited than CPUs to massively parallel workloads and offer, for the type of computation required by HLT1, a higher throughput per watt (and per Swiss franc) [40].

Each DAQ period — a data-taking period (RUN) followed by a long shutdown (LS), or roughly 5 to 6 years — brings about 200 new servers. The outgoing DAQ servers do not leave the system, however: they are redeployed to the HLT2 farm, where their operation continues until their repair potential is exhausted. The total lifetime adopted is approximately 10 years, reflecting the on-site maintenance and cascading practice of the online team. Like the other estimated lifespans in the model, it is propagated stochastically in the uncertainty analysis — a triangular distribution with a mode of 10 years and bounds of 6 and 12 years, reflecting shorter first deployments on one side and the limits of the cascade practice on the other (section 2.3.4). On top of this main flow, each new cycle brings additional purchases of servers destined directly for HLT2, which are difficult to quantify and are not modelled here.

Maintenance follows this cascading approach: very few components are purchased as spare parts. Failed servers are repaired without new components when possible, or by harvesting parts from decommissioned machines kept for this purpose. Once this stock is exhausted, the chassis enter the waste electrical and electronic equipment (WEEE) channel, following French treatment practices (see section 2.3.2). This cannibalisation practice extends the effective lifetime of the fleet.

None of the manufacturers involved publishes a detailed bill of materials or an environmental product declaration for the installed server models. The embodied impacts of the three server classes were therefore modelled exclusively in ResilioDB 2026.7, using the technical specifications of the configurations deployed at the online data centre. ResilioDB was selected because its parameterised models account for physical characteristics that strongly influence the manufacturing impacts of electronic components, whereas generic ecoinvent datasets do not provide sufficient technological resolution.

ResilioDB models CPUs and GPUs from their physical characteristics, including die area (mm^2), lithography node (nm) and thermal design power (TDP). The model covers die fabrication, including production losses, packaging and upstream transport. When the die area of a processor is not documented, ResilioDB estimates it using a random-forest model trained on approximately 4,000

CPUs with known die areas and reporting a mean absolute error of about 8 %. This estimation is reflected in the input-data quality score assigned by the database.

Other server components, including power supplies, storage devices and memory, are assessed from their relevant physical parameters. According to the ResilioDB documentation, these models draw on scientific literature and on additional sources whose details are not disclosed. Their complete construction therefore cannot be independently reproduced from the publicly available documentation.

The custom TELL40 readout cards hosted in the Class A servers are not represented in ResilioDB and are therefore excluded from the embodied-impact results.

Server end of life is modelled using the default ResilioDB treatment scenario, which assigns 54 % of the waste to recycling, 33 % to incineration and 13 % to landfill. Transport to treatment facilities and disassembly operations are not included.

Offline

Estimating the embodied impacts of the offline system requires a different approach from the one used for the online system. LHCb's offline computing resources are provided through the WLCG (section 2.3.1) and distributed across approximately forty sites worldwide, each with its own hardware fleet, replacement cycles and operating practices. The information required for a site-specific bottom-up inventory—including server models, component inventories and procurement dates—is not available at this scale. A top-down allocation approach is therefore used.

Rather than reconstructing the hardware fleet site by site, this approach applies an embodied-impact intensity expressed per logical core, or job slot, and per unit of time. The intensity is obtained by dividing the embodied impact of a reference server configuration by the number of logical cores it exposes to the batch system and by its calendar service life. It is then multiplied by the wall-clock occupation reported for LHCb in logical core-hours. The method therefore estimates the fraction of a reference server's lifetime capacity occupied by LHCb and provides an order-of-magnitude assessment consistent with the available accounting data.

The scope of the offline embodied-impact assessment is defined as follows:

- *Compute servers: included.* Their impacts are estimated using the logical-core-time allocation method described below.
- *Disk storage: included.* Data preservation is a central function of the grid. Disk impacts are modelled separately from computing using ResilioDB hard-drive models and the storage capacities provisioned for LHCb.
- *Tape storage: excluded.* The available information is insufficient to derive a defensible embodied-impact intensity covering cartridges, drives and tape libraries. Because tape stores a substantial share of the preserved data, this exclusion makes the offline embodied-impact estimate a lower bound and is discussed in Discussion.
- *Data centre infrastructure: excluded.* Buildings and technical installations, including cooling and electrical-distribution equipment, cannot be modelled consistently across the heterogeneous WLCG sites with the available data.
- *Data transfer infrastructure: excluded.* National and international research networks are shared among many users, and the available data do not support a defensible allocation of

their embodied impacts to LHCb.

In the absence of site-specific end-of-life data, the default ResilioDB end-of-life modelling described above is applied to the included servers and storage devices.

The grid servers are not dedicated to LHCb: each site shares its resources among several users. An allocation rule is therefore required to attribute only the share of embodied impacts associated with LHCb's use. The WLCG accounting introduced in section 2.3.1 provides two possible allocation metrics: the computing work delivered, expressed in HS23·h (figure 2.5), and the wall-clock occupation of computing resources, expressed in logical core-hours. The latter is retained here.

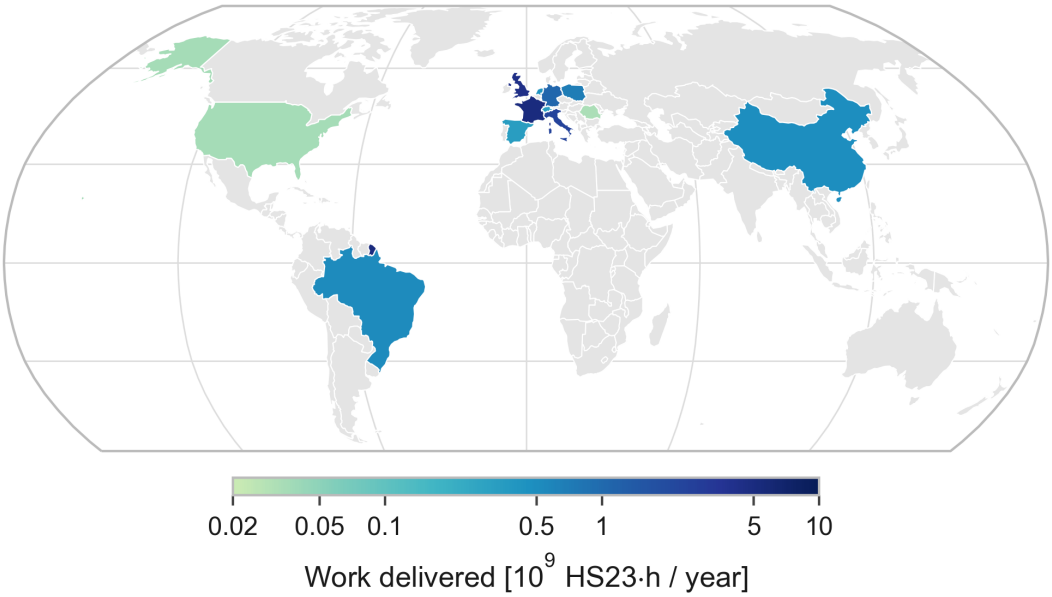


Figure 2.5: Computing work delivered to LHCb by the offline system in 2025 (CERN in Switzerland)

These quantities represent different aspects of computing activity. HS23·h measures the performance-normalised work delivered. Using it to allocate the embodied impact of a reference server would make the result depend on the performance difference between that server and the heterogeneous fleet that actually executed the workload. Wall-clock time instead records how long logical cores, or job slots, were allocated to LHCb. It therefore provides a direct measure of the share of schedulable server capacity occupied by the experiment.

The allocation is performed directly in logical core-hours, which is the unit reported by the WLCG accounting. Wall-clock time is retained rather than CPU time because a reserved job slot mobilises server capacity throughout the elapsed duration of the job, including periods during which the processor is waiting for data or other resources (figure 2.10).

The allocation follows the usual convention applied to capital goods in the LCA of IT: the embodied impact of a server is amortised over its full service life, and each user is attributed the fraction of its capacity that it occupies [18]. A reference server exposing n_{slot} logical cores provides, over a lifetime L, a total capacity of $n_{\text{slot}} \times L$ logical core-hours. The share attributed to LHCb is therefore:

$$E_{\text{emb,LHCb}} = E_{\text{emb,server}} \times \frac{T_{\text{LHCb}}}{n_{\text{slot}} \times L} = EI \times T_{\text{LHCb}}, \tag{2.3}$$

where $E_{\text{emb,server}}$ is the embodied impact of the reference server, n_{slot} the number of logical cores it exposes to the batch system, L its lifetime in hours and T_{LHCb} the wall-clock occupation attributed

to LHCb, expressed in logical core-hours. The corresponding embodied-impact intensity is:

$$EI = \frac{E_{\text{emb,server}}}{n_{\text{slot}} \times L}, \tag{2.4}$$

and is expressed in the relevant impact unit per logical core-hour, for example in grams of CO_2eq per logical core-hour for climate change.

The denominator represents the complete calendar capacity of the reference server rather than only the fraction of time during which its processors execute instructions. No utilisation factor is therefore introduced. This convention is consistent with the treatment of the online fleet, whose embodied impacts are also amortised over calendar time. Capacity that remains unoccupied is not attributed to LHCb; the resulting estimate is consequently lower than one based on the complete redistribution of idle capacity among users.

The lifespan of the servers and its uncertainty distribution follow the ones adopted for the online IT servers.

This formulation does not require the number of servers used by LHCb to be estimated. The recorded logical core-hours are converted directly into an equivalent fraction of the lifetime capacity of the reference server, irrespective of whether the workload was distributed over many small machines or fewer high-density systems.

It remains to select the reference server used to calculate EI. Because the hardware fleet operated by the WLCG sites is not inventoried in sufficient detail, the calculation relies on the configuration-specific ResilioDB estimates established for the online server classes (section 2.3.2). Class B is retained as the reference for the central offline estimate. It is a recent CPU-only configuration and therefore resembles the type of hardware used for the predominantly CPU-based reconstruction and simulation workloads of the grid.

Class A is not retained as a possible reference. It was designed for the GPU-accelerated HLT1 workload, and its hardware configuration is not representative of a conventional grid compute server. Moreover, its HS23 score characterises only its CPU performance and does not account for the computing capacity provided by its GPUs. The custom TELL40 cards are not included in its ResilioDB model, as stated above.

Class C is used only as an alternative reference in the sensitivity analysis. It represents the older and heterogeneous CPU hardware that still constitutes a substantial part of the online fleet. The WLCG accounting gives an average performance of approximately 14.8 HS23 per job slot for the hardware that executed LHCb workloads in 2025. This value lies between the surveyed values of class B (15.66 HS23 per job slot) and class C (10.94 HS23 per job slot). These performance values are used only to assess the representativeness of the reference classes; HS23·h does not enter the embodied-impact calculation itself.

Under the central ten-year lifetime assumption, the climate-change intensity of class B is:

$$EI_B = \frac{2{,}267\,\text{kg CO}_2\text{eq}}{256 \times 10 \times 8{,}760\,\text{h}} = 0.101\,\text{g CO}_2\text{eq per logical core-hour.} \tag{2.5}$$

For class C, the corresponding intensity is:

$$EI_C = \frac{1{,}994\,\text{kg CO}_2\text{eq}}{198 \times 10 \times 8{,}760\,\text{h}} = 0.115\,\text{g CO}_2\text{eq per logical core-hour.} \tag{2.6}$$

The class C intensity is therefore approximately 13.7 % higher than that of class B. The same allocation is applied separately to every EF 3.1 impact category available from the ResilioDB models. Class B defines the central result, while class C quantifies the sensitivity to the use of a different fleet (section 2.3.4).

The embodied impacts of disk storage follow the same amortisation principle but use capacity rather than computing occupation as the allocation unit. The number of drives attributed to LHCb is obtained by dividing the average provisioned capacity by the capacity of a reference drive. A 16 TB drive is retained centrally; the available ResilioDB models for 12, 14, 18, 20 and 22 TB drives are used as sensitivity variants, with the number of drives adjusted in each case to maintain the same total capacity. Their embodied impacts are amortised using the same IT-equipment lifespan distribution as the servers.

Storage is allocated according to the capacity provisioned for LHCb, because this capacity is reserved for the experiment whether or not every byte is occupied. The annual mean provisioned capacity is used centrally, while the lower occupied capacity is evaluated as an alternative allocation scenario.

The main uncertainties of this approach are the representativeness of class B for the heterogeneous WLCG fleet, the assumed IT-equipment lifetime and the unknown capacity distribution of the installed drives. The reference-server and drive-capacity choices are examined through sensitivity scenarios, while lifetime is propagated stochastically. Tape storage and grid-site infrastructure remain unquantified, the offline embodied-impact results are then to be considered as a lower bound.

2.3.3 Operational impacts

On-site

The on-site operational assessment covers the electricity and water consumed by the LHCb online data centre. Electricity supplies both the IT equipment and the auxiliary services, principally cooling. Water is used mainly for the adiabatic assistance of the cooling system.

Staff commuting and other staff travel are excluded from the system boundary. Their allocation would require poorly constrained information on transport modes, teleworking, assignment sites and the purpose of individual journeys. Component replacements, such as air filters, are instead included in the embodied impacts.

The IT equipment is cooled by indirect free-air cooling with adiabatic assistance. The facility comprises six prefabricated IT modules, each served by four air handling units (AHUs), for a total of 24 units. In each AHU, the air from the IT rooms circulates in a closed loop, while a separate outside-air stream passes along the opposite side of an air-to-air heat exchanger. Heat is transferred between the two streams without bringing outside air into direct contact with the electronic equipment, thereby limiting the ingress of humidity, dust and pollutants. When dry heat exchange is insufficient, water is sprayed on the outside-air stream which lowers the air temperature and provides additional cooling capacity.

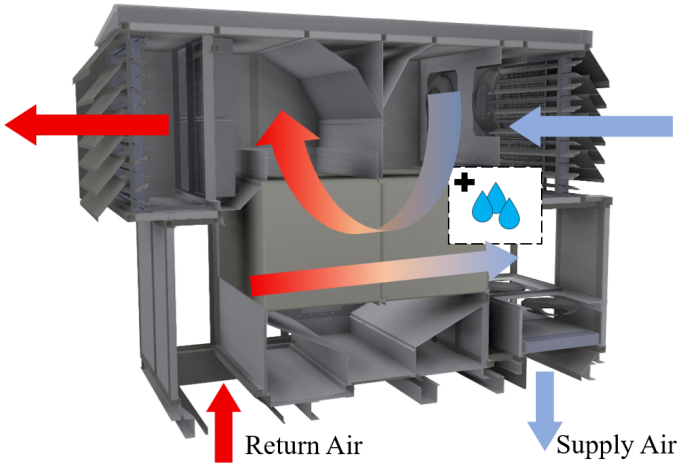


Figure 2.6: Schematic of an air handling unit

During operation, the air returning from the servers' hot aisles is cooled in the AHUs before being recirculated towards the server inlets. The target inlet temperature is approximately 27 °C [13, 41]. The cooling capacity of the air-to-air exchanger depends principally on the temperature difference between the two air streams, their flow rates and the performance of the exchanger.

As the outside-air temperature increases, the temperature difference available for heat exchange decreases. The control system can respond by increasing the airflow through the exchanger, which raises fan electricity consumption, or by activating the adiabatic assistance. In the latter case, water is sprayed into the outside-air stream, and its evaporation lowers the air temperature before it reaches the exchanger. The operation of the cooling system can therefore involve a trade-off between fan electricity and water consumption. The available measurements do not, however, establish whether a given activation of the adiabatic system was necessary to maintain the temperature setpoint or was selected to limit fan electricity consumption.

The Power Usage Effectiveness (PUE) is defined as the ratio between the total energy consumed by the data centre and the energy delivered to the IT equipment [16]:

$$\mathrm{PUE} = \frac{E_{\mathrm{total}}}{E_{\mathrm{IT}}} = \frac{E_{\mathrm{IT}} + E_{\mathrm{cooling}} + E_{\mathrm{losses}} + E_{\mathrm{aux}}}{E_{\mathrm{IT}}} \geq 1, \tag{2.7}$$

where E_{IT} is the energy consumed by the IT equipment, E_{cooling} that of the cooling system, E_{losses} the electrical distribution losses and E_{aux} the remaining auxiliary consumption. In accordance with ISO/IEC 30134-2 [16], E_{total} includes all energy crossing the data-centre boundary, including lighting, security, fire protection and building-management systems [16]. An ideal PUE equals 1.

The worldwide annual mean PUE was approximately 1.54 in 2025 (figure 2.7).

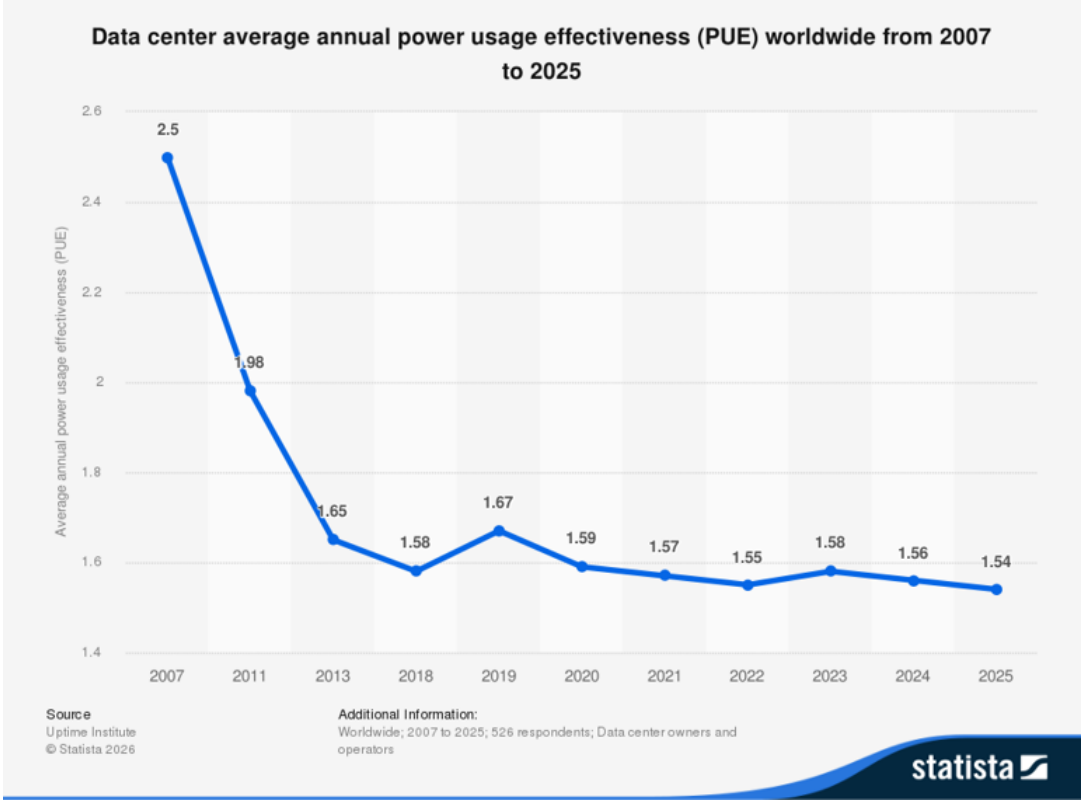


Figure 2.7: Evolution of the worldwide average annual data-centre PUE [11]

For the LHCb online data centre, the annual mean PUE measured in 2025 was 1.053 [42]. For context, CERN reports a design target of approximately 1.1 for its Prévessin data centre and a PUE of approximately 1.5 for Meyrin; the value reported for the WLCG in 2023 was 1.45 [22, 43]. The low value measured at Point 8 is consistent with the limited auxiliary energy required by indirect free-air cooling under the climatic conditions encountered during the year.

Both PUE and WUE depend partly on the meteorological conditions of the year considered. The values reported here should therefore be interpreted as observations for the 2025 reference year rather than as climate-independent characteristics of the facility. The relationship between outside temperature and the measured use of adiabatic cooling is examined directly using the Point 8 monitoring data in figure 2.8.

Water efficiency is quantified using the Water Usage Effectiveness (WUE) [17], defined as the ratio between annual water consumption and the energy delivered to the IT equipment:

$$\mathrm{WUE} = \frac{V_{\mathrm{water}}}{E_{\mathrm{IT}}} \quad [\mathrm{L\,h/kW}]\,, \tag{2.8}$$

where V_{water} is the annual volume recorded for the adiabatic cooling system. The 2025 measurements show a clear seasonal association between daily water consumption and outside temperature (figure 2.8). Consumption is concentrated between April and October and becomes nearly zero during winter. Daily consumption exceeded 1 $\mathrm{m^3}$ on 212 days and 10 $\mathrm{m^3}$ on 107 days.

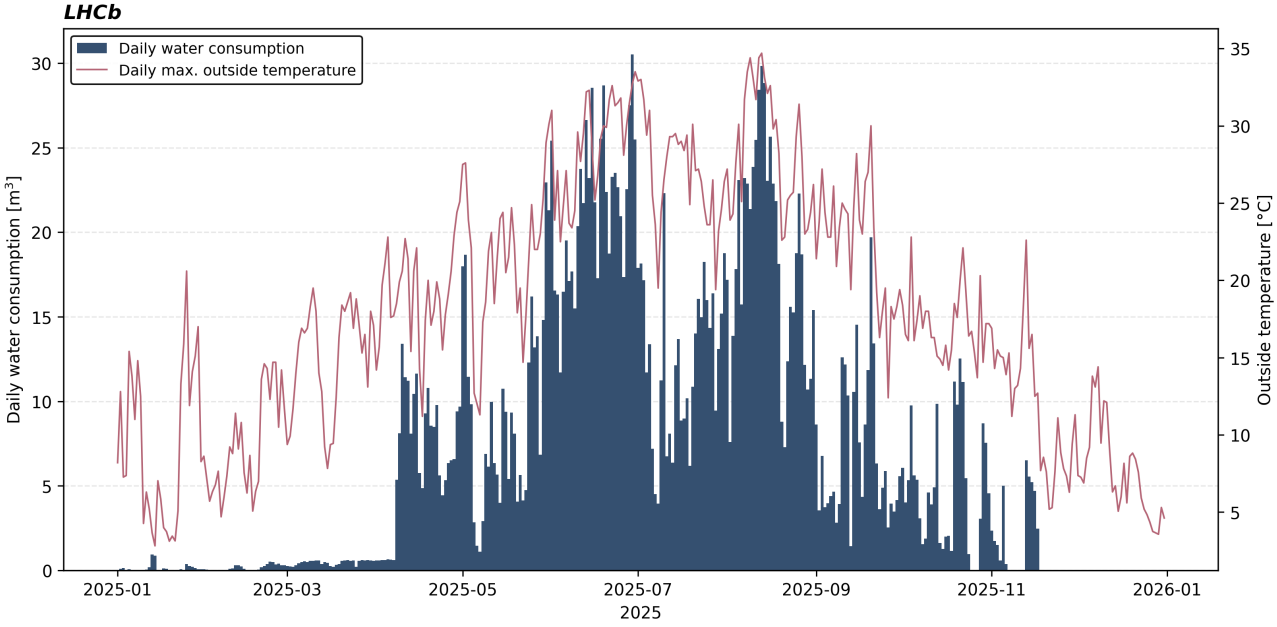


Figure 2.8: Daily water consumption of the adiabatic cooling system and daily maximum outside temperature at Point 8 in 2025

Integrated over 2025, the adiabatic cooling system consumed approximately 2,516 m^3 of water. More than 95 % of this volume was consumed between April and October, with a monthly maximum of approximately 650 m^3 in June. Applying equation (2.8) (over a full year then) gives a WUE of approximately 0.22 L h/kW.

The environmental impacts associated with electricity consumption are computed, for every impact category j, as:

$$I_{\text{elec},j} = E_{\text{site}} \times f_{\text{elec},j}, \tag{2.9}$$

where E_{site} is the total metered electricity consumption and $f_{\text{elec},j}$ is the corresponding electricity impact factor. For climate change, this factor is a carbon intensity expressed in kg CO_2eq/kWh; for the other EF v3.1 categories, it is the characterised impact per unit of electricity.

Two electricity-data sources are retained for complementary purposes. Ecoinvent 3.12 provides a consistent profile across all sixteen EF v3.1 impact categories, but its country electricity mixes mainly represent 2022 and therefore do not match the studied year. Electricity Maps provides annual carbon intensities for 2025 with a life-cycle boundary, but only for the climate-change indicator. The dataset used here is the Electricity Maps 2025 yearly dataset redistributed in `@tgwf/co2` v0.18.0 [28, 29].

The complete EF v3.1 multi-impact profile is consequently calculated using the ecoinvent country mixes for both the online and offline systems. It includes the impacts of the metered water consumed by the online system. The headline climate-change result and the online–offline comparison use the Electricity Maps 2025 annual averages on both sides and cover electricity-related operational impacts only. Direct online water use is excluded from this headline because an equivalent contribution cannot be quantified for the offline sites; this keeps the comparison of the two electricity supplies on the same boundary. The difference between the ecoinvent and Electricity Maps climate results is reported as a sensitivity scenario in section 2.3.4. Other carbon-intensity sources

examined later, including Ember and CaDI, are used for inter-source comparison.

Average rather than marginal emission factors are used. Marginal factors describe the response of the electricity system to a change in demand and therefore belong to a consequential perspective, whereas the present inventory follows an attributional approach. Their use would consequently be inappropriate for the allocation of the site's observed annual consumption [44, 45].

The calculation uses average electricity mixes rather than residual mixes. As a general rule, each WLCG site is assigned the average mix of its host country. CERN-PROD, the Tier 0 site, constitutes an exception: although it is geographically located in Switzerland, CERN is supplied by the French grid, and its electricity consumption is therefore assigned the French mix [6]. Given the large share represented by CERN Tier 0, assigning the Swiss electricity mix would substantially misrepresent the resulting impacts. The same French mix is used for the LHCb online data centre, which is physically located in France at Point 8.

Offline

The offline operational assessment covers the electricity attributed to LHCb computing and disk storage across the WLCG, including the electricity used by site infrastructure through the PUE. For every EF v3.1 impact category, these electricity impacts are characterised using the corresponding ecoinvent country mix for each site. Climate change is additionally calculated using the Electricity Maps annual averages for 2025.

The quantified perimeter therefore comprises two components: the electricity associated with computing activity and that associated with disk storage. The latter is estimated separately below. Tape storage is not quantified because the available data do not support a consistent estimate of the electricity consumed. Inter-site data transfers are also excluded following the screening assessment presented at the end of this subsection.

No direct electricity measurements are available for the complete set of WLCG sites serving LHCb. Compute electricity must therefore be estimated from the activity reported in the WLCG accounting (section 2.3.1). Each occupied logical core is assigned an average power representing its share of the complete server, including the processor, memory, power supplies and internal fans. Electricity used by cooling, electrical distribution and other site infrastructure is then included through the PUE. The compute energy attributed to LHCb at site s is:

$$E_s^{\text{compute}} = T_s \times P_{\text{core}} \times \text{PUE}_s, \tag{2.10}$$

where T_s is the wall-clock occupation reported for LHCb in logical core-hours, P_{core} is the average complete-server power attributed to one occupied logical core, and PUE_s accounts for the additional electricity consumed by the data-centre infrastructure.

Due to the lack of data, a uniform value of $\text{PUE}_s = 1.45$, taken from the holistic assessment of WLCG energy needs, is applied to all sites [22].

For impact category j, the impacts associated with compute electricity are then calculated as:

$$I_{\text{offline},j}^{\text{compute}} = \sum_s E_s^{\text{compute}} \times f_{j,c(s)}, \tag{2.11}$$

where $f_{j,c(s)}$ is the electricity impact factor assigned to site s. The index $c(s)$ refers to the country whose grid supplies the site rather than necessarily to its geographical location. It generally corresponds to the host country; the specific treatment of CERN, which is assigned the French electricity mix, is described below. For the headline climate-change result, $f_{j,c(s)}$ is the Electricity Maps carbon intensity expressed in kg CO_2eq/kWh; for the complete EF v3.1 profile, it is the corresponding ecoinvent impact factor.

The value of P_{core} was estimated from an unpublished dataset gathering the technical specifications of 212 processor models, supplied by N. Szczepanek and D. Giordano (CERN) [21].

Processor TDP cannot be used directly as P_{core}: it is a processor-level design rating rather than a measurement of the electrical power drawn by a complete server. It is therefore used here only as a proxy. For each processor model i, the complete-server power per logical core is estimated as:

$$p_i = \frac{\text{TDP}_i}{N_{\text{threads},i}} \times u_{v(i)}, \tag{2.12}$$

where $N_{\text{threads},i}$ is the number of hardware threads exposed by the processor and $u_{v(i)}$ is vendor-specific uplift factor. This factor approximates the contribution of the rest of the server, including memory, the motherboard, internal fans and power-supply losses. Cooling and the remaining facility-level consumption are excluded because they are accounted for separately through the PUE.

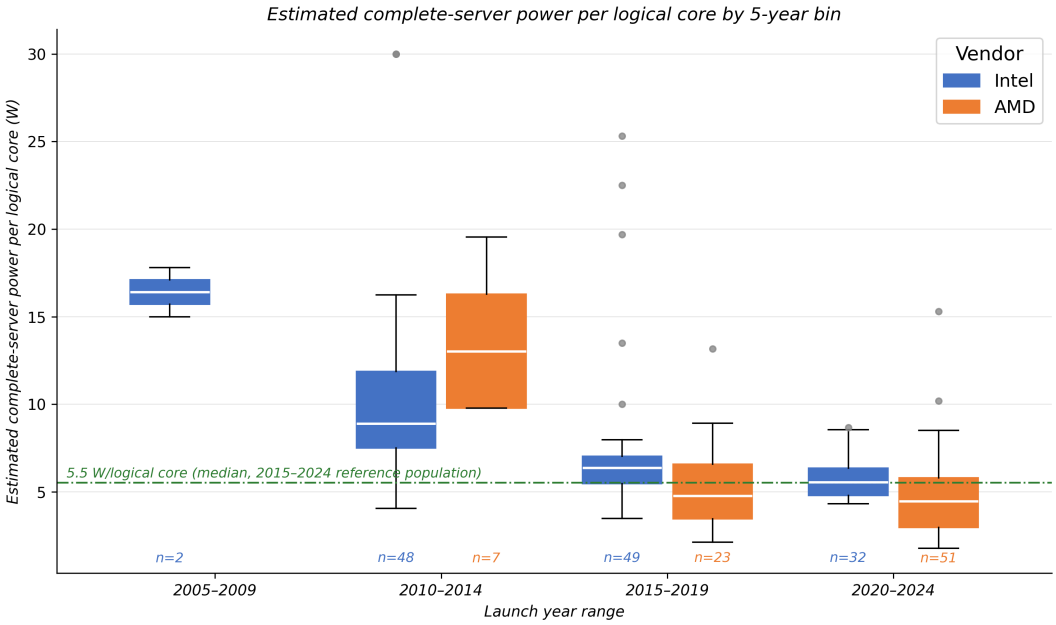


Figure 2.9: Estimated complete-server power per logical core, calculated from processor TDP, hardware-thread count and a confidential vendor-specific uplift factor

Figure 2.9 shows a general decrease across successive processor generations. To define a central proxy for the hardware operating in 2025, the reference population is restricted to the 155 processor models launched between 2015 and 2024. This cut-off is a modelling hypothesis rather than an inventory of the processors actually installed at WLCG sites.

The median of this population is retained as the central value:

$$P_{\text{core}} = 5.5\,\text{W} \quad \text{per occupied logical core.} \tag{2.13}$$

The 10th and 90th percentiles, 2.9 W and 7.9 W per logical core, delimit the distribution assigned to P_{core} in the uncertainty analysis (section 2.3.4). The parameter remains a fleet-level proxy applied uniformly to all sites, not a measurement of the servers that actually executed the LHCb jobs.

Of the two accounting metrics defined in section 2.3.1, WCT is retained as T_s in equation (2.10) because it records the full duration for which logical cores were allocated to LHCb jobs, including time spent waiting for data, input/output or other resources. CPU time, by contrast, records active instruction execution. The observed CPU efficiency of approximately 96 % implies a difference of about 4 % between the two metrics (section 2.3.1), which is small compared with the uncertainties assigned to P_{core} and the PUE. No separate CPU-time scenario is therefore retained.

This allocation does not assume that server power is proportional to CPU utilisation. Servers continue to draw power below full load, as illustrated in figure 2.10 [46, 47].

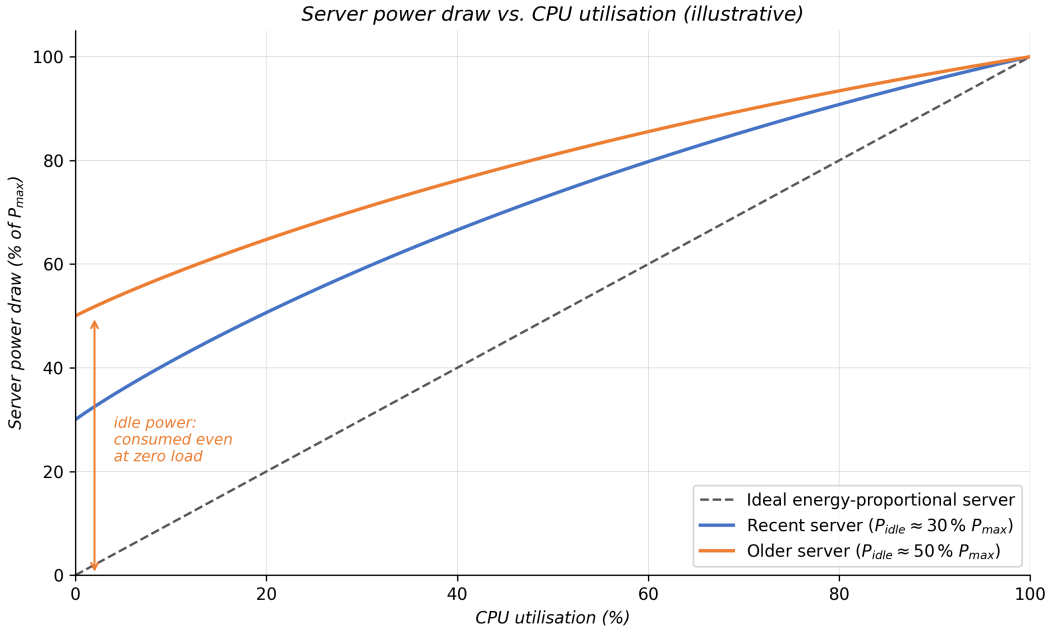


Figure 2.10: Illustrative server power as a function of CPU utilisation, normalised to peak power P_{max}. The curves follow the empirical model of Fan, Weber, and Barroso [47], $P(u) = P_{\text{idle}} + (P_{\text{max}} - P_{\text{idle}})(2u - u^{1.4})$, for two assumed idle-power fractions. The figure illustrates non-proportional behaviour and is not used to derive P_{core}

The conversion of electricity into environmental impacts follows the same attributional, annual-average approach as the on-site assessment (section 2.3.3). For climate change, each site is assigned the 2025 Electricity Maps carbon intensity of its host country. CERN-PROD is the only exception: its electricity is supplied by the French grid and is therefore assigned the French value despite its Swiss location [48]. The versioned Electricity Maps snapshot used in the calculation is distributed with the open-source CO2.js package `@tgwf/co2` v0.18.0 [28, 29]. For the complete EF v3.1 profile, the corresponding ecoinvent 3.12 country mixes are used. Alternative carbon-intensity sources are evaluated in the sensitivity analysis (section 2.3.4).

The electricity used by disk storage is estimated from the capacity provisioned for LHCb at each site in the WLCG accounting. Because the number and capacity of the installed drives are not

published, this capacity is multiplied by a specific power $w = P_{\text{HDD}}/S_{\text{HDD}}$, where P_{HDD} is the power of an enterprise hard-disk drive and S_{HDD} its capacity. Manufacturer specifications for Seagate Exos X16 and X18 drives give values of approximately 0.3 to 0.6 W/TB, depending on the drive capacity and operating state [49, 50]. A central drive power of 6.5 W and a representative capacity of 16 TB are adopted, giving $w = 0.41\,\text{W/TB}$.

Unlike computing, provisioned storage remains powered continuously. The energy is therefore integrated over the calendar months to account for changes in capacity during the year:

$$E_s^{\text{disk}} = \sum_m C_{s,m} \frac{P_{\text{HDD}}}{S_{\text{HDD}}} h_m \times \text{PUE} \tag{2.14}$$

where $C_{s,m}$ is the disk capacity provisioned for LHCb at site s in month m and h_m is the number of hours in that month. The compute estimate covers the worker-node hardware, whereas this term represents the bulk-storage capacity reported separately by the WLCG; the two contributions are therefore treated as distinct. No additional RAID or erasure-coding overhead is applied because the physical redundancy configuration of each site is unavailable.

Inter-site data transfers over LHCOPN/LHCONE are excluded following an order-of-magnitude screening. Integrating the hourly LHCb throughput exported for 2025 gives 41.8 PB. Applying the fixed-network core-and-access intensity of 0.0016 kW h/GB reported by Europe Economics [51] gives approximately 66.9 MW h, equivalent to 0.31 % of the 21.51 GW h already included for online and offline electricity [52]. Although this factor is not a direct measurement of LHCOPN/LHCONE, the result supports excluding transfers at the precision of this study.

Water use is not assessed for the offline system. Cooling technologies and water monitoring differ between grid sites, and available meters generally cover computing centres shared among several users. Neither the site-level cooling data nor the allocation keys needed to attribute water consumption to LHCb are available.

2.3.4 Uncertainty analysis

The purpose of the uncertainty analysis is to determine how the quantified result changes when uncertain inputs and modelling relationships vary. Continuous uncertainties are propagated by Monte Carlo (MC) simulation, whereas discrete methodological choices are evaluated as scenarios.

Applying a common uncertainty method is challenging because the result combines three computational branches. openLCA and ecoinvent provide exchange distributions, ResilioDB provides point estimates accompanied by quality scores, and the offline model uses simple equations whose main parameters are estimated from accounting data and proxy hardware. These sources are divided into uncertainty groups.

Uncertainty groups and distributions

Four groups are used for stochastic propagation:

1. *Dataset quality (G1)* covers the uncertainty carried by the source datasets for their declared scope. In the sampled openLCA product systems, native ecoinvent exchange distributions are used as published. For ResilioDB components, this group is represented by a factor derived from the *manufacture source* score.
2. *Representativeness (G2)* covers the additional uncertainty introduced when a dataset or mea-

sured proxy represents different technology, geography or time. It includes the representativeness part of the openLCA pedigree assessment, the ResilioDB *environmental assessment quality* factor, the grid-server power per logical core and the disk specific power.

3. *Input data (G3)* covers the reliability and completeness of data collected for this study. It includes the corresponding pedigree factors, the ResilioDB *input data quality* factor, the grid PUE and the uncertainty on the class C fleet size.

4. *Lifespans (G4)* are kept separate because they affect every annualised embodied impact through the inverse relationship in equation (2.2).

Table 2.4 summarises the implementation of these groups. The complete list of stochastic parameters and bounds is reported in table 4.

Table 2.4: Partition of the stochastic uncertainties used in the model

Group	**Sources propagated**
G1 — Dataset quality	Native ecoinvent background distributions in the three sampled openLCA chains; ResilioDB factor mapped from the *manufacture source* score
G2 — Representativeness	Foreground temporal, geographical and technological correlations; ResilioDB *environmental assessment quality*; P_{core} and disk specific power
G3 — Input data	Foreground reliability and completeness; ResilioDB *input data quality*; class C fleet size and grid PUE
G4 — Lifespans	Shared lifespan draws for IT, electromechanical and long-lived equipment; deterministic replacement interval for air filters

For foreground data assessed with the ecoinvent Data Quality System (DQS), geometric variances are additive in log-space [53]:

$$\ln^2\text{GSD}_{\text{tot}} = \ln^2\text{GSD}_{\text{basic}} + \sum_k \ln^2\text{GSD}_k. \tag{2.15}$$

The five pedigree indicators can therefore be separated into a data factor, based on reliability and completeness, and a representativeness factor, based on temporal, geographical and technological correlation:

$$a_i = \bar{a}_i \times f_{\text{data},i} \times f_{\text{repr},i}, \tag{2.16}$$

where $\bar{a}_i$ is the amount encoded in the inventory. Both factors are independent lognormal variables with a median of one.

Indicators marked as not assessed are treated as neutral and add no variance.

The lifespan distributions are based on the lifespan assumptions introduced in section 2.3.2. IT hardware, including online and offline servers and storage, follows a triangular distribution $(6; 10; 12)$ years. Air-handling units, inverters and pumps follow $(15; 20; 25)$ years. Air filters retain their deterministic six-month replacement interval. Components whose deterministic lifespan was capped at 50 years follow the one-sided triangular distribution $(50; 50; 100)$ years, because the cap indicates that their service life is expected to exceed 50 years.

ResilioDB does not provide probability distributions. Its three quality scores are therefore mapped onto the corresponding DQS indicators and converted into lognormal factors with a median of one [53, 54]. This provides a reasonable approximation of the uncertainties. Because the scores are inspired by the PEF data-quality framework, their conversion inherits the limitations of qualitative-to-quantitative mappings [33, 55].

Propagation in the three model branches

The openLCA background distributions are propagated with its native MC engine through the *olca-ipc* Python interface[1]. Corrected ecoinvent 3.12 files and openLCA 2.6.2 were used after issues affecting pedigree uncertainty and linked upstream results had been identified; the provenance controls and recomputation procedure are documented in Provenance of the openLCA uncertainty chains.

Three openLCA chains were retained after screening the contribution of each product system: the on-site operational block and the offline operational block were each sampled over 500 native iterations, and the container over 200. The other eleven infrastructure components remain in the balance with their deterministic background results and sampled Python factors and lifespans, but their native ecoinvent background variance is not propagated. Among the infrastructure components, only the container's native background uncertainty was sampled. Propagating uncertainty in its foreground inputs would have required three additional openLCA chains to separate data-quality from representativeness effects. Because the container contributes 0.46 % of the ecoinvent climate-change result, these chains were not run. Its central foreground inventory and lifespan uncertainty remain included.

The chain lengths were selected for estimating variances rather than only central values. Under a normal distribution approximation, the relative standard error of the sample variance is $\sqrt{2/(n-1)}$; this gives a 90 % confidence half-width of approximately 10 % at $n = 500$ and 16 % at $n = 200$. Because the actual chains are non-normal and some have heavy tails, this expression is used only as a theoretical benchmark. A pilot chain and bootstrap diagnostics on the simulated draws provided the empirical check of the retained sample sizes.

The openLCA product systems are calculated per unit. Quantities, foreground factors and annualisation by the sampled lifespans are then applied in Python. The native exchange correlations induced by shared upstream processes are preserved within each product-system chain. Separate product-system chains are paired independently, so covariance of their background results is not captured.

At each MC iteration, the nominal impact of each ResilioDB component is scaled by three log-normal factors derived from its quality scores: manufacture-source quality (G1), environmental-assessment representativeness (G2) and input-data quality (G3). The factors associated with the class B server are shared between its online contribution and its use as the offline reference model because both calculations use the same ResilioDB estimate.

The analytical offline model is propagated directly in Python. Its operational calculation samples three parameters:

- P_{core}, triangular $(2.9; 5.5; 7.9)$ W per occupied logical core;
- the grid PUE, triangular $(1.40; 1.45; 1.50)$;
- disk specific power, triangular $(0.30; 0.406; 0.60)$ W/TB.

Wall-clock time is kept deterministic because it is an accounting record. Carbon intensities are also kept deterministic within each run because their providers publish no variance; differences between sources are assessed through scenarios.

[1] `https://github.com/GreenDelta/olca-ipc.py`

Offline embodied impacts follow the allocation in section 2.3.2. The reference-server impact is divided by its number of logical cores and its sampled lifetime, then multiplied by LHCb's logical core-hours. Disk storage is represented by the number of drives required for the provisioned capacity and uses the same IT-lifespan draws. The server-class choice and HDD capacity are treated as scenarios and therefore do not contribute to the variance groups. No RAID or erasure-coding overhead is added, so the grid-storage embodied impact remains a lower bound.

Assembly and variance decomposition

At each iteration, five blocks are assembled: online embodied impacts from openLCA, online embodied impacts from ResilioDB, on-site operational impacts, offline operational impacts and offline embodied impacts.

The offline operational multi-impact chain is rescaled by the sampled ratio $(P_{\text{core}}/\bar{P}_{\text{core}}) \times (\text{PUE}/\overline{\text{PUE}})$. Disk electricity is added separately and is rescaled by PUE and disk specific power, but not by P_{core}. The complete EF 3.1 profile uses the ecoinvent electricity chains and includes the climate-change impact of on-site water use. The Electricity Maps headline is calculated analytically from electricity consumption and country carbon intensities for both systems. Direct online water is excluded from this headline because direct water use cannot be quantified for the offline sites; the comparison consequently covers electricity-related climate impacts on the same basis for both systems.

The openLCA simulations provide 200 or 500 native iterations, whereas the final Python assembly uses 50,000 draws to propagate the remaining parameters. At each assembly iteration, one existing result is selected at random from each openLCA chain; the same result may be selected more than once. All sixteen impact-category results from the selected openLCA iteration are kept together, thereby preserving their joint variation. This procedure allows the native chains to be combined with the 50,000 Python draws, but it does not create additional ecoinvent information.

For the group decomposition, the assembly is repeated with one uncertainty group active and the other three fixed at their nominal values. The variance of each restricted assembly is divided by the variance of the full assembly. These first-order shares do not include multiplicative interactions; the difference between their sum and unity is reported as an interaction residual.

Confidence intervals on the group shares are obtained by resampling the *native* openLCA iterations and propagating them through the complete assembly. The reported 90 % intervals use 200 bootstrap replications. Bootstrapping the 50,000 assembled draws directly would instead measure the precision of repeated resampling from the same few hundred native outputs.

Scenario analysis and reporting

Five discrete choices are evaluated one at a time by re-running the complete stochastic assembly with the same random numbers. Each scenario therefore produces a distribution and can affect its location and width. Table 2.5 lists the retained axes.

The comparison between Electricity Maps and ecoinvent does not only include a temporal bias: the two sources also differ in method and system boundary. It is therefore interpreted as the sensitivity of the climate result to the electricity modelling basis. It indicates that the multi-impact profile depends on the ecoinvent electricity model, but it does not quantify a dating bias for the other fifteen categories.

Table 2.5: Sensitivity scenarios. Central values are those used for the main results

Parameter	Central	Variants	Comment
Offline reference class	B (Supermicro)	C (generic)	Alternative proxy for the heterogeneous grid fleet (section 2.3.2)
HDD reference capacity	16 TB	12, 14, 18, 20, 22 TB	Unknown capacity distribution of the installed drives
Storage allocation	Provisioned capacity	Occupied capacity	Reserved capacity compared with bytes actually stored
Carbon-intensity dataset	Electricity Maps	Ember, CaDI	Differences between carbon-intensity providers
Electricity modelling basis	Electricity Maps (2025)	ecoinvent 3.12 mixes	Combined sensitivity to reference year, modelling method and system boundary

For each impact category, the MC outputs are summarised by the mean, standard deviation, median and 5th/95th percentiles. The deterministic calculation at the central parameter values remains the reference balance for descriptive results. The arithmetic MC mean is reported as the expectation under the assigned distributions; its difference from the deterministic result reflects both the inverse-lifespan relationship and the right-skewed lognormal factors whose median, rather than mean, equals one.

Ionising radiation has the heaviest right tail and the least stable upper percentile of the sixteen categories. All draws are retained, but the indicator is reported as a secondary result with its convergence diagnostics. Maintenance and replacement parts absent from the inventories, local electricity generation at grid sites and uncertainty in EF 3.1 characterisation factors are not quantified. Together with processes excluded from the boundary, they are discussed qualitatively in Discussion.

3. Results

This chapter presents the results for the processing of LHCb data in 2025. At the central parameter values and using Electricity Maps annual averages for operational electricity, the climate-change impact of the complete system is 3,100 t CO_2eq, of which 963 t is attributed to the online branch and 2,137 t to the offline branch.

Only the *operational* contributions depend on the electricity-data source; embodied impacts are unchanged between the ecoinvent and Electricity Maps results. The first part of this chapter presents the complete multi-impact profile calculated with ecoinvent. The second focuses on climate change and uses Electricity Maps for operational electricity. The definitions and units of the EF v3.1 impact categories are provided in the appendix (table 5).

3.1 Online system

3.1.1 Embodied impacts

The sixteen environmental indicators in figure 3.1 share two main hotspots: servers and storage. Their relative importance varies across impact categories. The infrastructure and network have lower impacts. For the infrastructure, this is mainly due to its long lifespan, which reduces its annualised impacts below those of the two main contributors (tables 7 and 8).

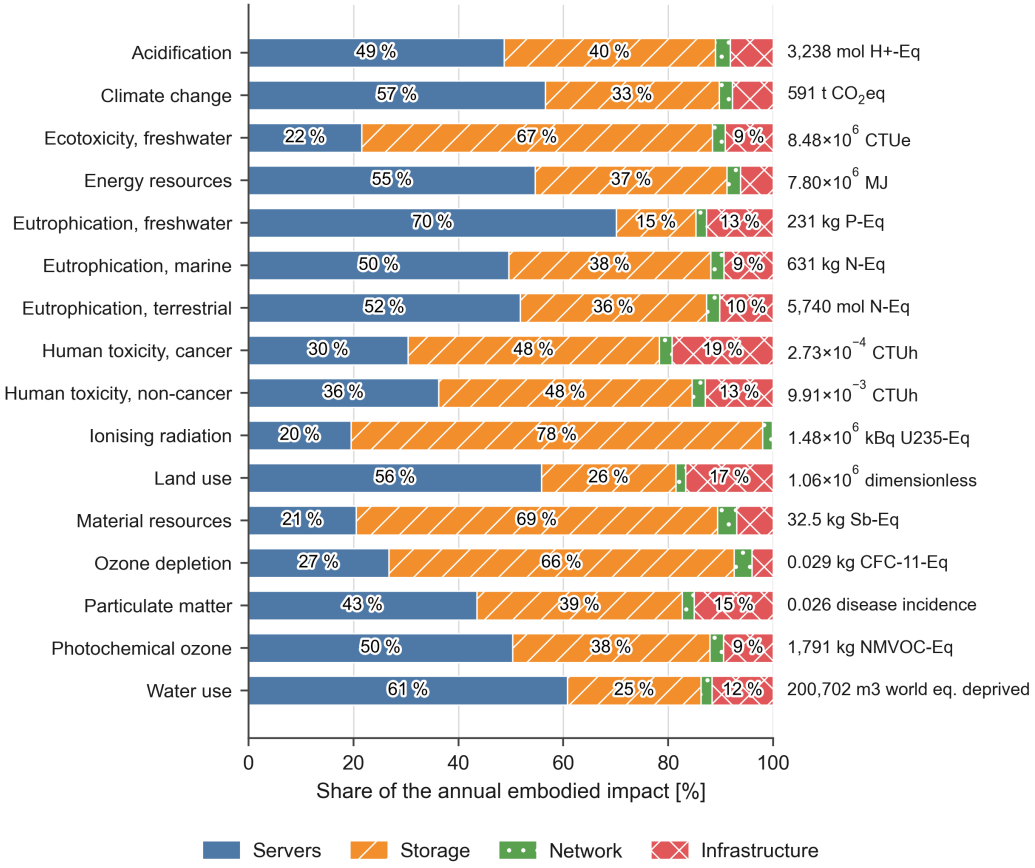


Figure 3.1: Annual embodied impacts of the online system by equipment category — EF v3.1

For climate change, most of the impact comes from servers, and manufacturing dominates across all

IT equipment (see figure 4 in the appendix). Across the sixteen categories, manufacturing accounts for 91.9 % of the climate-change impact. End of life reaches 35.6 % of freshwater ecotoxicity, 13.1 % of marine eutrophication and 9.9 % of human toxicity, non-cancer (figure 6).

Figures 3.2 and 3.3 present the climate-change impact of one unit from each server class, first by component and then by manufacturing process. The three classes have similar profiles, with RAM as the main component contributor.

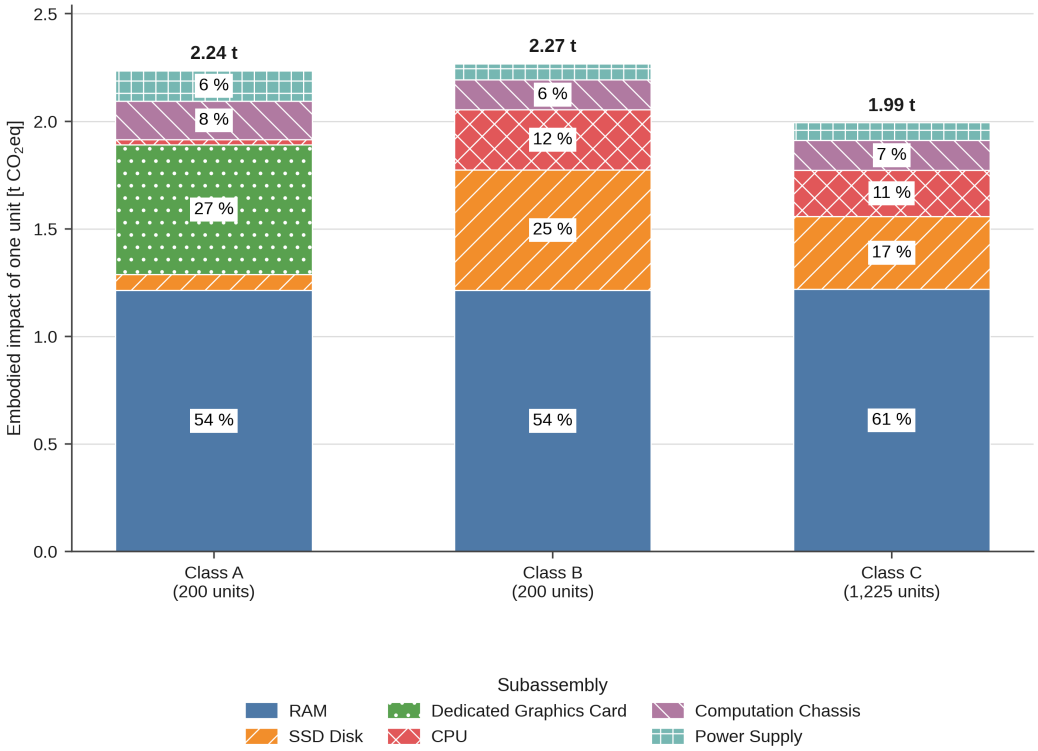


Figure 3.2: Embodied climate-change impact of one unit from each server class, by physical subassembly

By manufacturing process, the impact is dominated by the *logic wafer*. However, the processor impacts differ among the three classes: 11.4 kg CO_2eq for an EPYC 7502 chip, 26.7 kg for a Xeon E5-2680 v4 and 139.9 kg for a Xeon Platinum 8592+. This difference persists when the processor impacts are expressed per physical core: the respective values are 0.36, 1.91 and 2.19 kg CO_2eq per core. The core count therefore cannot explain the higher impact of the Xeon Platinum 8592+; the manufacturing parameters used by ResilioDB must also be considered (figure 5).

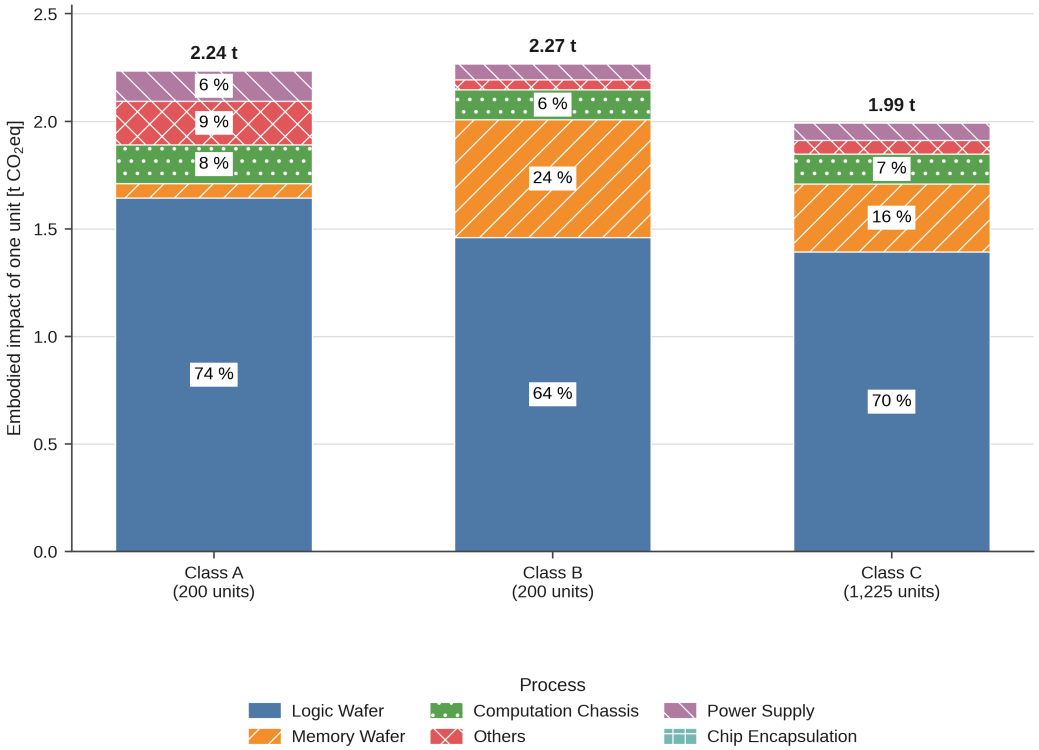


Figure 3.3: Embodied climate-change impact of one unit from each server class, by manufacturing process

Logic wafer manufacturing is the main contributor. One of the parameters used by ResilioDB to estimate a CPU impact is its estimated die area. For the Xeon 8592+, this area is three times larger than the maximum value recommended by the database for maintaining good result quality. Its high impact appears to result from a larger die area than that of the other processors — likely related to its 64 physical cores, compared with 32 and 14 — combined with poor model representativeness. This limitation is reflected in the processor quality score and then in the server score.

3.1.2 Overall online impacts

Operational impacts dominate fourteen of the sixteen categories. Embodied impacts dominate only mineral resources and freshwater ecotoxicity (figure 3.4).

Overall, the online profile combines operational dominance across most categories with a substantial embodied contribution from servers and storage. The low carbon intensity of the French electricity supply limits the climate contribution of use, which increases the relative importance of hardware manufacture and lifespan in the online result. The implications for procurement and renewal are discussed in Discussion.

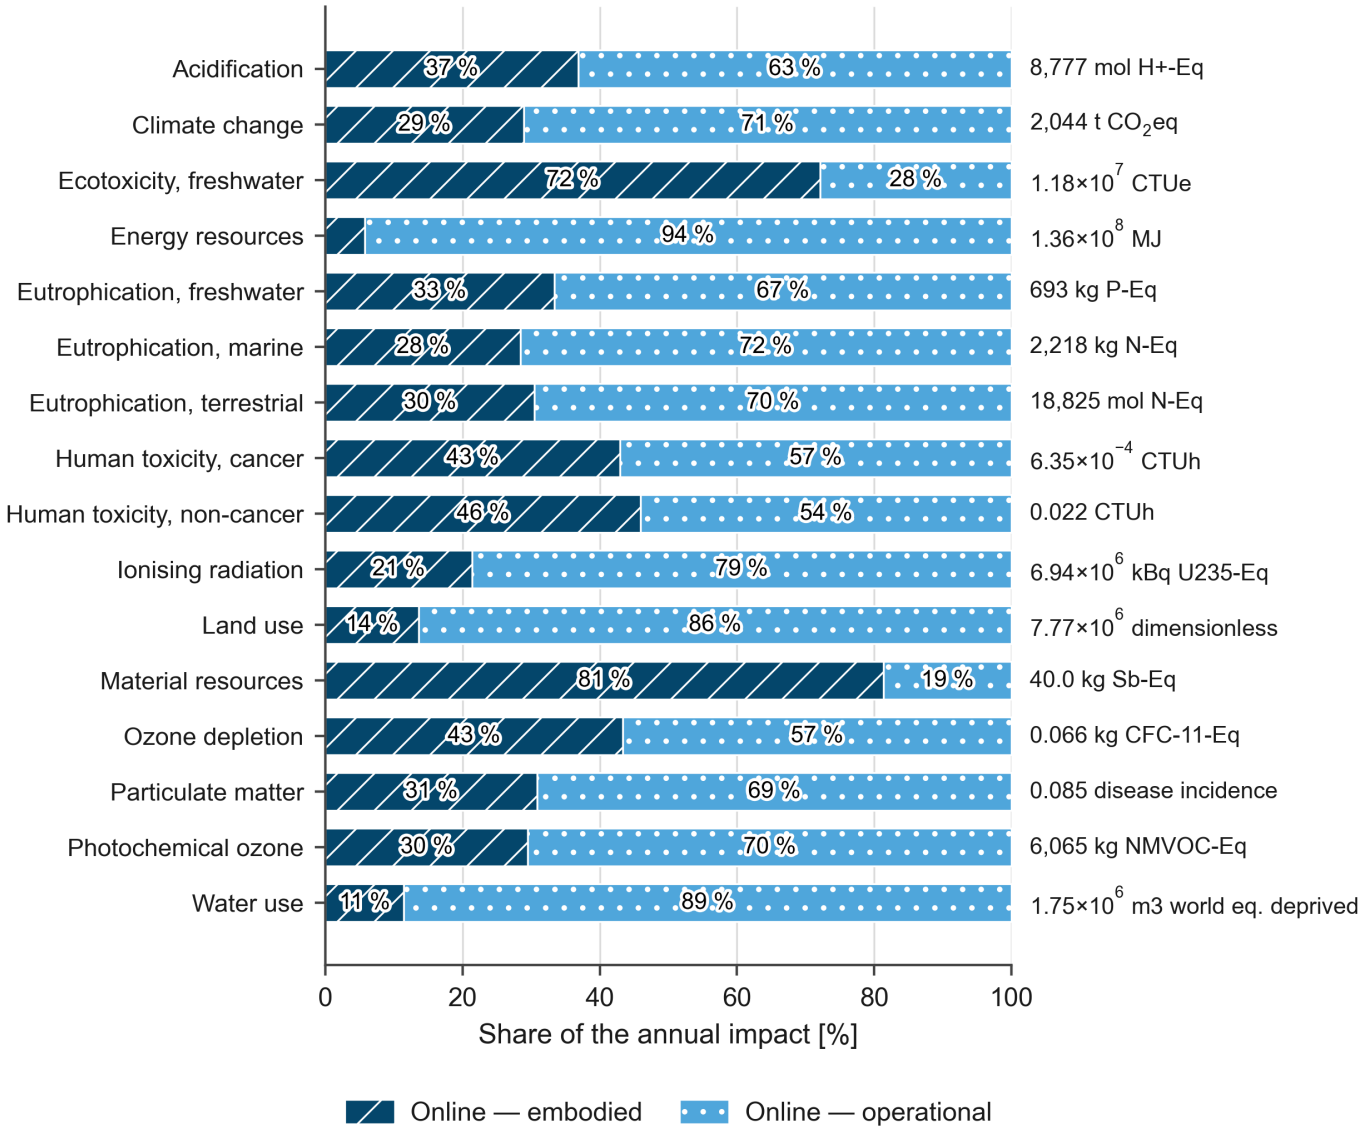


Figure 3.4: Total annual impacts of the online system — EF v3.1

3.2 Offline system

For consistency with the preceding analysis, the Tier-0 contribution from CERN ($\neq$ LHCb) is assigned to France in the offline results because the electricity consumed is supplied by the French grid, keeping a location based approach for such a big contributor would bring unnecessary wrong values.

3.2.1 Embodied impacts

Storage contributes more than servers to the embodied impacts. Compared with the online results, this difference is explained by the large number of offline disks: 9,144 units in the configuration based on 16 TB drives. The complete embodied-impact profile is provided in the appendix (figure 7).

3.2.2 Overall offline impacts

Water consumption at offline data centres is not included because neither the data nor a sufficiently robust estimation method were available. The "water use" category in the complete profile provided in the appendix (figure 8) therefore includes only upstream water use in the electricity and equipment supply chains.

Operational impacts dominate most categories, except freshwater ecotoxicity and mineral resources. As in the online system, these two categories are driven by embodied impacts. For the offline system, the main source is hard-drive production, especially the drive itself and its power supply (see the Sankey diagram in the appendix, figure 10).

3.3 Complete system: multi-impact profile

The total system impacts are presented in figure 3.5. Numerical values, uncertainty intervals and additional figures are provided in the appendix (tables 16 and 17).

Electricity consumed by the grid is the main hotspot for most impact categories. Ionising radiation is an exception: its largest contribution is online electricity, due to the high share of nuclear power in the French electricity mix. Mineral resources are dominated by offline embodied impacts, reflecting the metals contained in the approximately 9,000 disks provisioned for the experiment. Freshwater ecotoxicity follows the same pattern for the same reason.

To provide a common order-of-magnitude scale across categories, figure 3.6 expresses the characterised impacts in person equivalents using the EF 3.1 normalisation factors [34]. Normalisation is used only to support interpretation and does not replace the absolute category results.

The normalised profile highlights four categories: fossil-resource use, freshwater eutrophication, ionising radiation and mineral-and-metal resource use. They are retained in the multi-impact screening but are not discussed at the same level of detail as climate change, which is the principal interpretation focus defined in section 2.3.1.

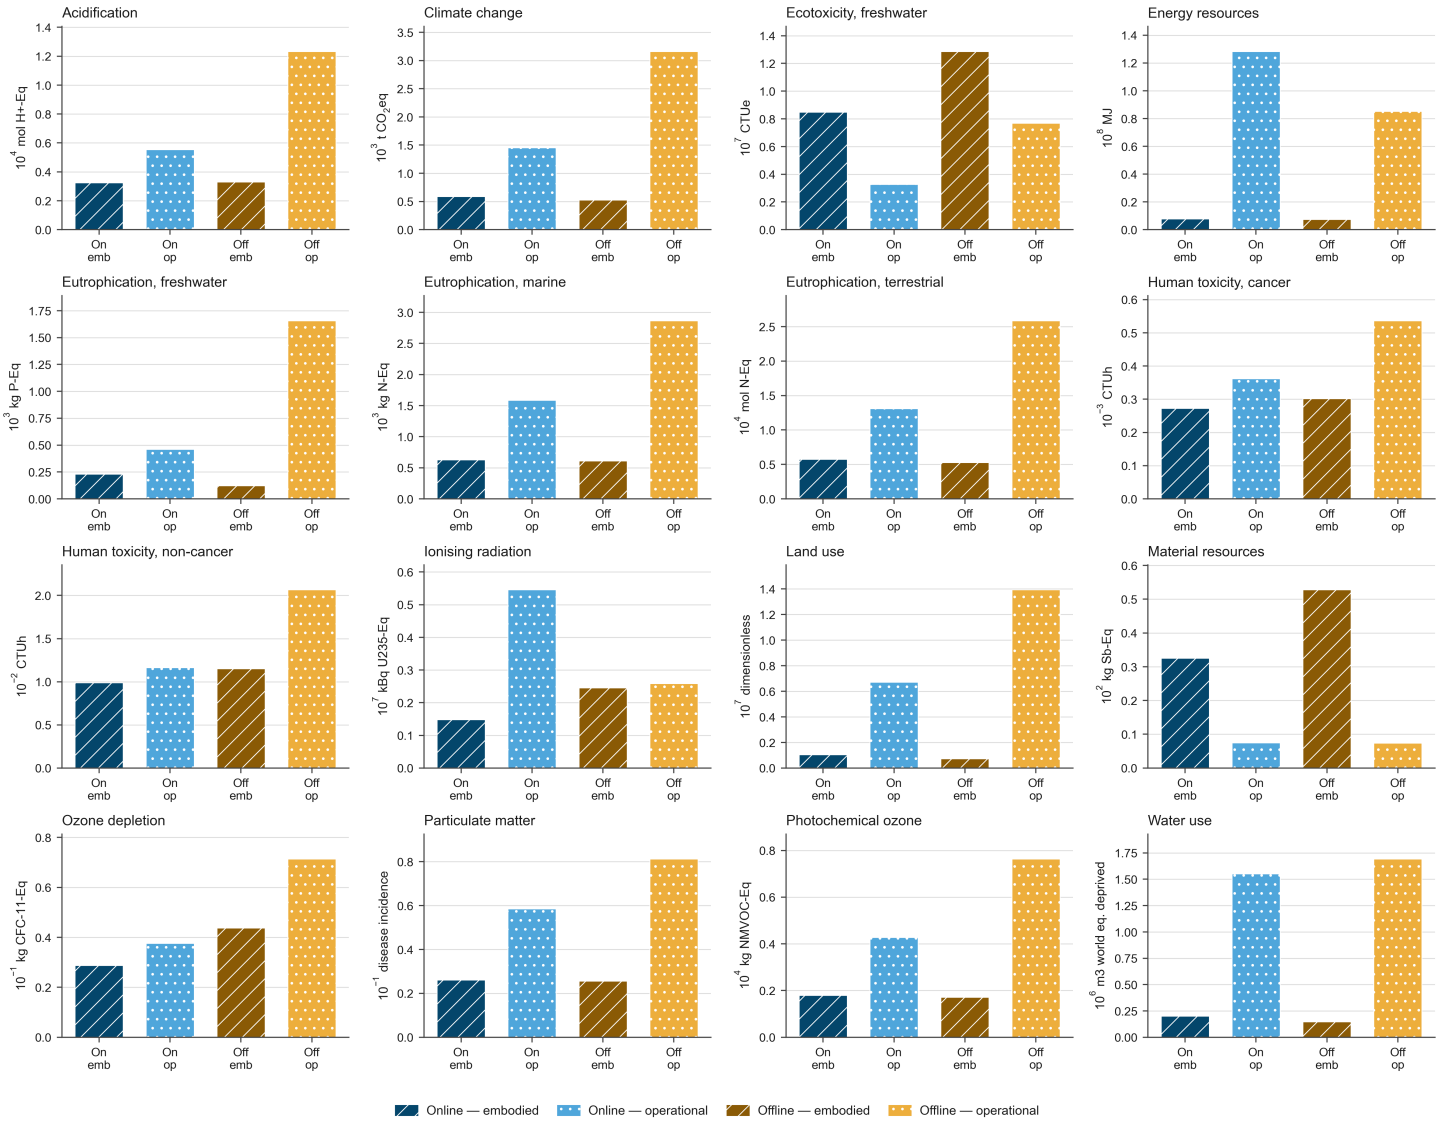


Figure 3.5: Total annual impacts of the complete system at the central parameter values — EF v3.1

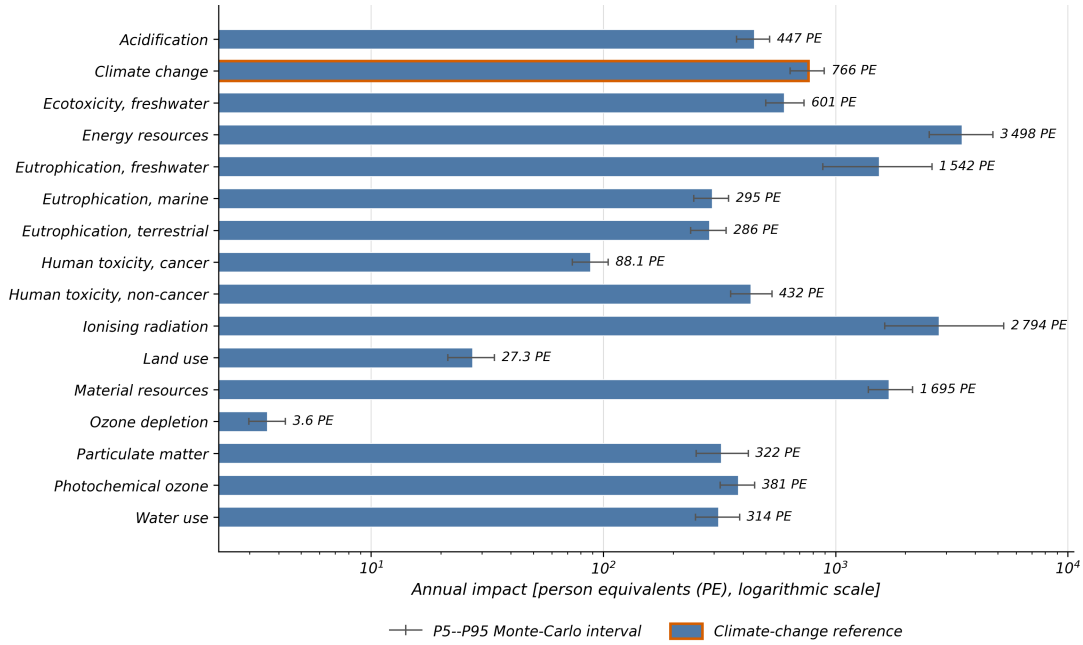


Figure 3.6: Person-equivalent normalisation of the complete-system EF 3.1 results for 2025, using the EF 3.1 normalisation factors

3.4 Online and offline climate-change impacts

At the central parameter values, the complete-system result is 3,100 t CO_2eq with Electricity Maps and 5,734 t with ecoinvent (figure 3.7). The two electricity-data sources differ in reference year, modelling method and system boundary; their difference therefore represents the combined sensitivity to the electricity-modelling basis.

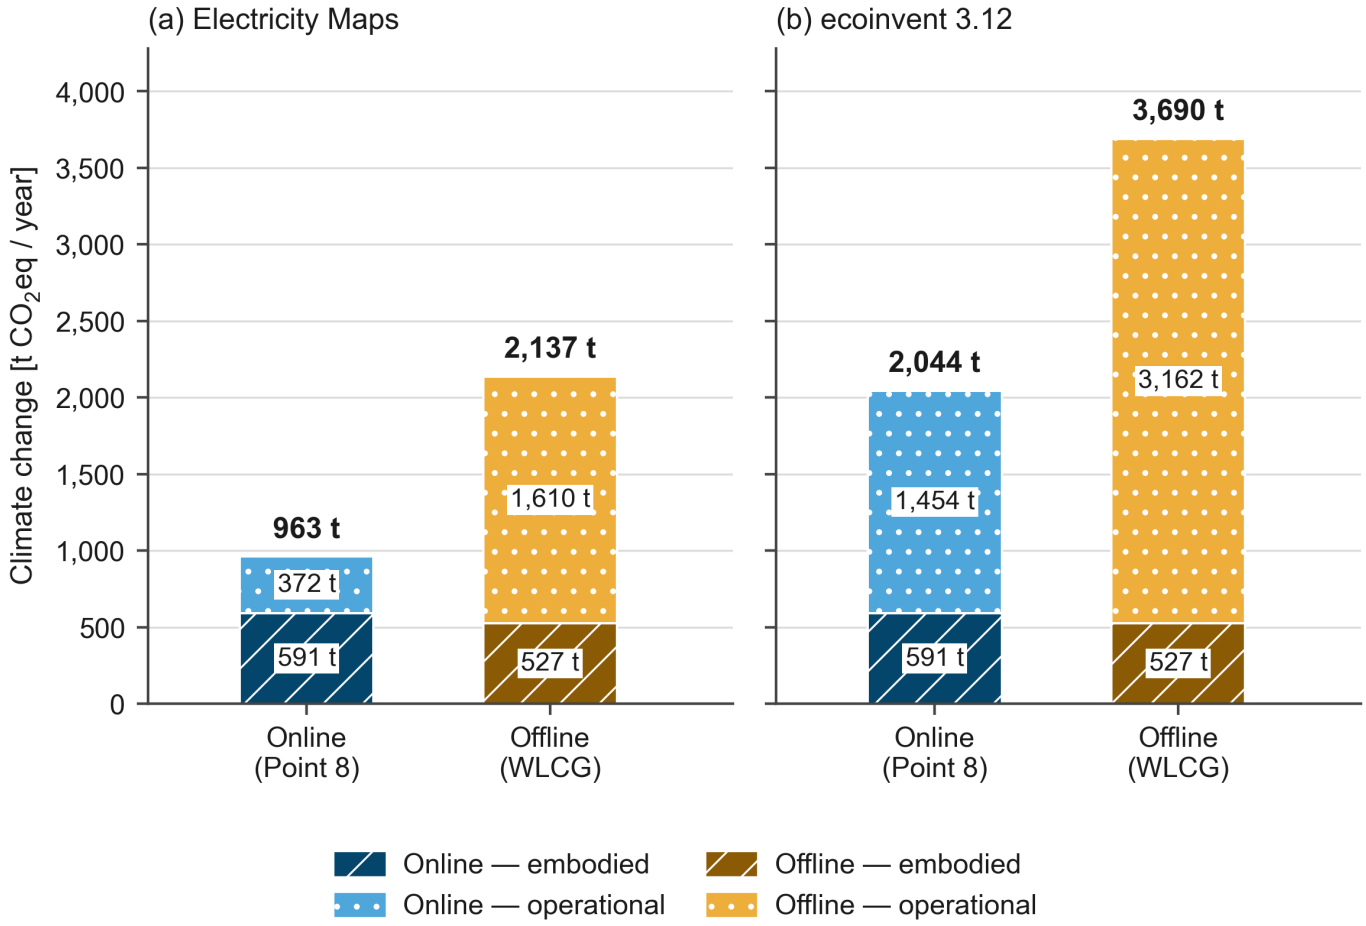


Figure 3.7: Deterministic climate-change balance at the central parameter values: (a) Electricity Maps; (b) ecoinvent 3.12

The electricity-data source changes both the magnitude of the result and the ratio between the online and offline operational contributions.

At the central parameter values, embodied impacts account for 591 t CO_2eq of the 963 t online total, whereas operational electricity accounts for 1,610 t of the 2,137 t offline total. The detailed online and offline breakdowns are provided in the appendix (figures 12 and 13). The online electricity factor is 31 g CO_2eq/kWh, compared with a weighted average of 184 g CO_2eq/kWh for the offline sites. The complete balance by contribution is provided in the appendix (table 18).

Most of the carbon impact is generated in Europe (figure 3.8). This geographical distribution must be compared with the work delivered.

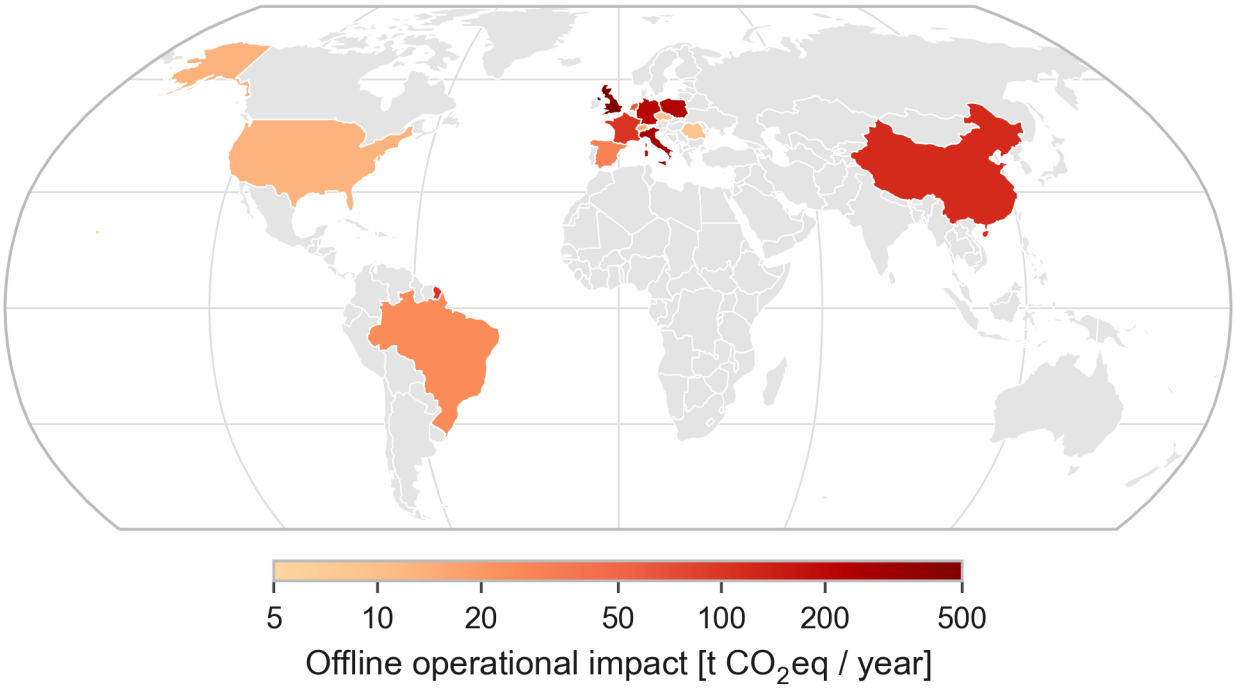


Figure 3.8: Operational climate-change impact of LHCb offline computing in 2025 by host country — Electricity Maps

This comparison is provided in the appendix (figure 14). France delivers 31.9 % of the grid computing work while accounting for 6.4 % of its emissions. Poland delivers 4.4 % of the work and accounts for 15.3 % of the emissions. The ratio of emission share to work share differs by a factor of 17 between the two countries.

Additional maps, including computing work delivered in HS23·hours by country and the electricity mixes of host countries, are provided in the appendix (figures 15, 16 and 2.5).

3.5 Uncertainty analysis results

3.5.1 Monte Carlo simulation

Figure 3.9 presents the distribution of the total climate-change impact. Across the 50,000 draws, the mean is 3,182 t CO_2eq and the P5–P95 interval ranges from 2,644 to 3,729 t. The result obtained with the central parameters (3,100 t) is 2.7 % below the mean. This difference is expected because annualised impacts depend on the inverse of lifespans and the lognormal factors preserve the nominal value as their median rather than their arithmetic mean. The deterministic result remains the reference balance for the descriptive figures; the Monte Carlo mean and percentiles characterise how that balance changes under the assigned uncertainty distributions (table 21).

Figure 3.10 retains the electricity-source comparison on the Monte Carlo basis. The bar subdivisions are MC means and the whiskers are the P5–P95 intervals of the corresponding system totals.

The distributions differ between the two systems. The offline distribution is nearly symmetric, whereas the online distribution is skewed; the detailed distributions are provided in the appendix (figure 17).

This difference follows from the structure of their variance. For the online system, electricity consumption is metered and its carbon intensity is not sampled, so all variance comes from embodied impacts. For the offline system, operational impacts account for 89 % of the variance. Their un-

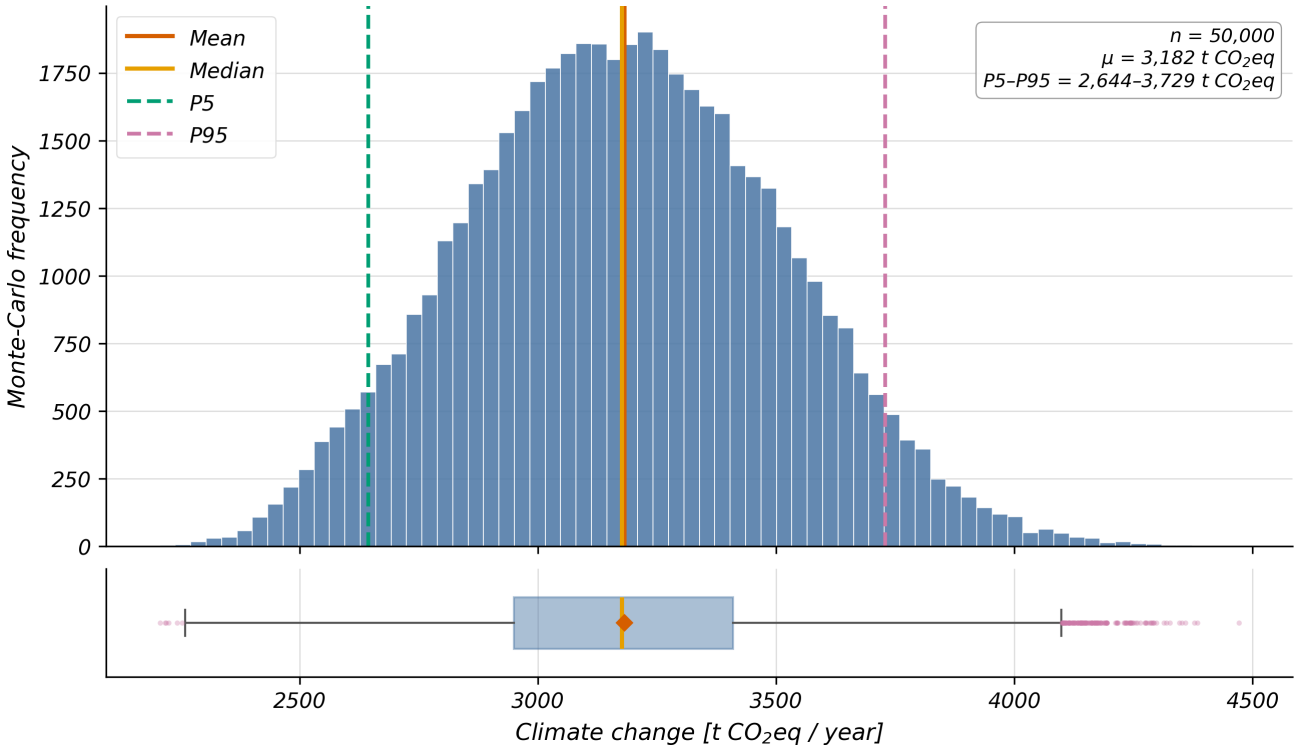


Figure 3.9: Distribution of the annual climate-change impact of the complete system in 2025 — Electricity Maps annual averages

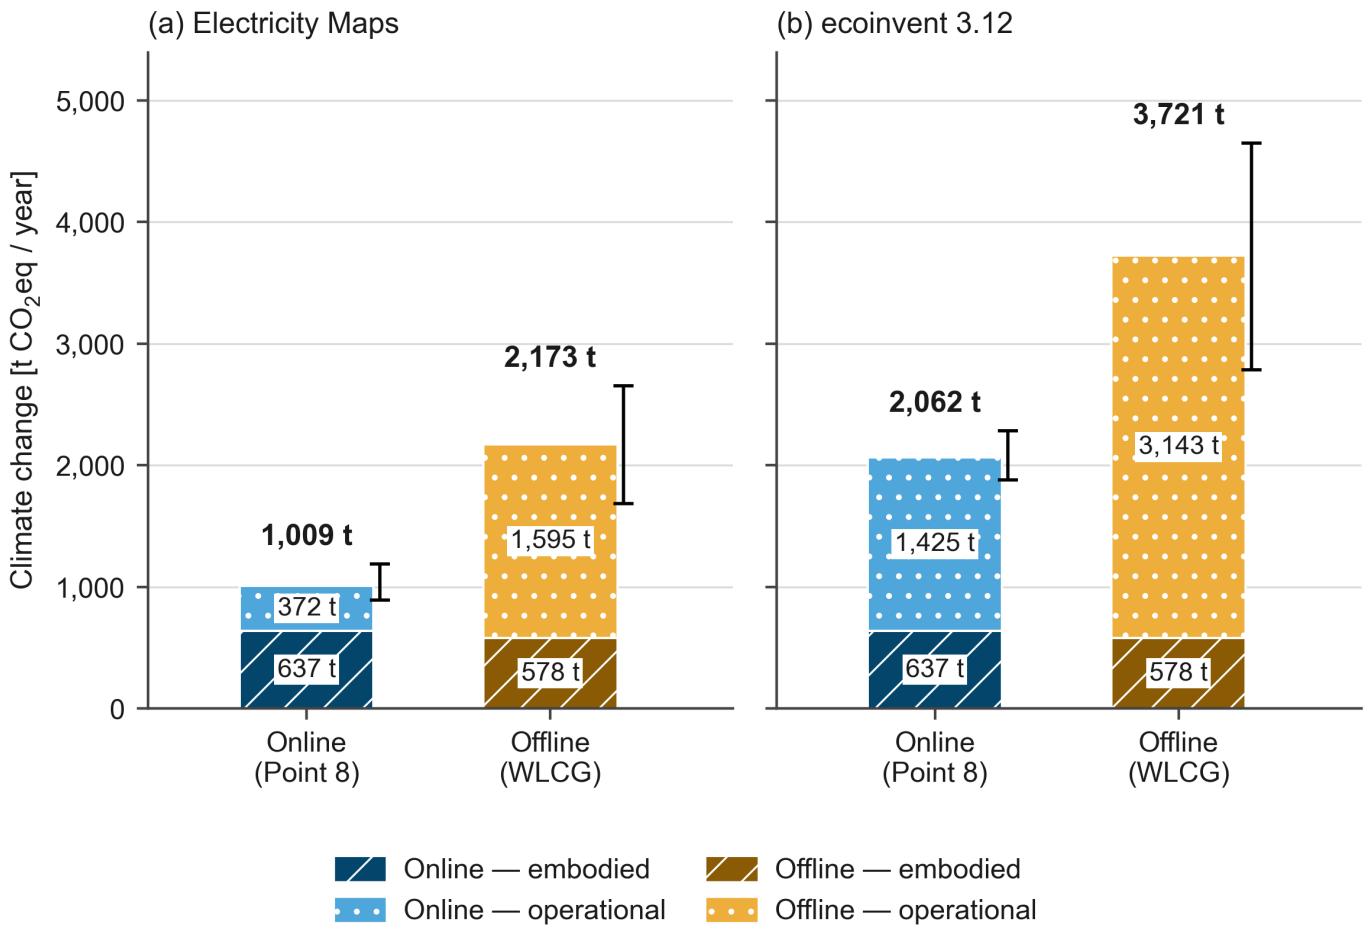


Figure 3.10: Monte Carlo climate-change balance by electricity-data source: means by contribution and P5–P95 intervals of the online and offline totals

certainty depends mainly on the product $P_{\text{core}} \times \text{PUE}$.

The distributions of the other fifteen environmental indicators are provided in the appendix (figure 19 and table 27).

3.5.2 Contributions to the variance

Representativeness is the main source of uncertainty in both cases (figure 3.11). It explains 70.5 % of the variance of the main indicator and 85.7 % of that of the climate-change row in the multi-impact profile. Lifespans rank second, while input data contribute little to the variance. The residual below 1 % indicates that interactions among groups are negligible and that their effects are nearly additive.

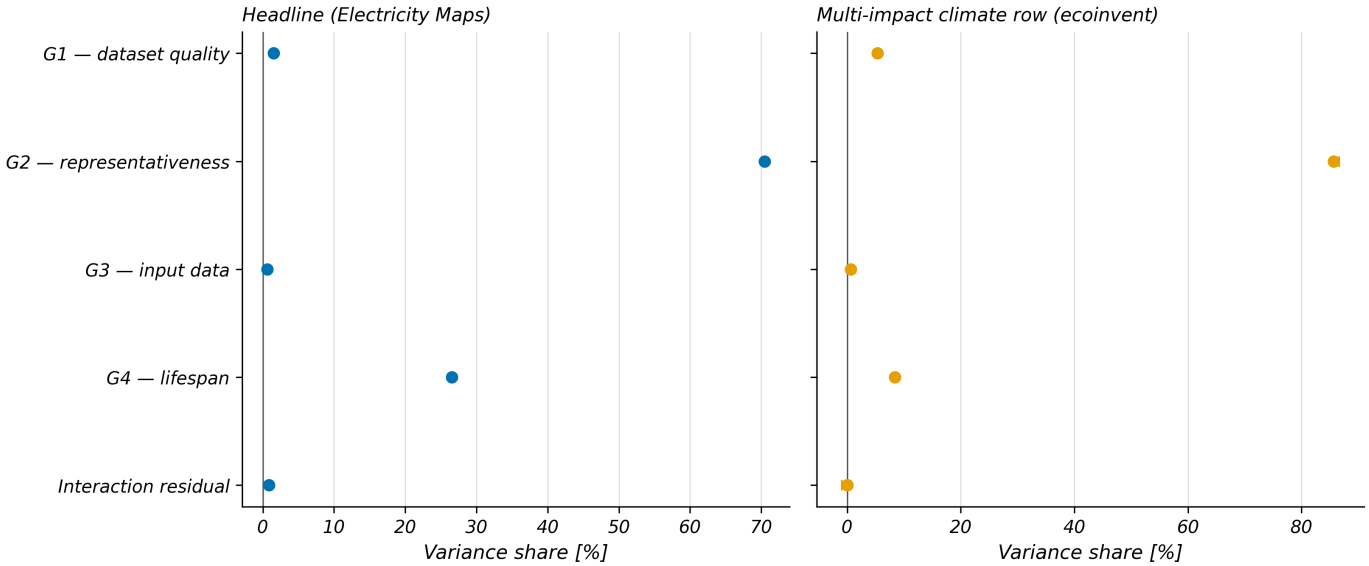


Figure 3.11: First-order contributions of the four uncertainty groups to the variance of the climate-change result

The predominance of representativeness reflects possible differences between the datasets and the actual fleet, especially the average power per logical core, disk specific power and the representativeness factors assigned to proxy equipment models. The choice of class B or C as the offline reference server is evaluated separately as a discrete scenario.

The electricity model affects the contributions of the four groups. With ecoinvent, the dispersion of background electricity data increases the contribution associated with dataset quality and reduces the relative contribution of lifespans. The group contributions therefore depend on the electricity basis used to calculate the climate-change impact.

3.5.3 Scenario variants

The preceding analysis propagates continuous parameters, whereas this section examines five discrete modelling choices. For each variant, one choice is changed and the complete Monte Carlo simulation is recalculated using the same random numbers. The difference shown in figure 3.12 is the change in the mean relative to the central scenario.

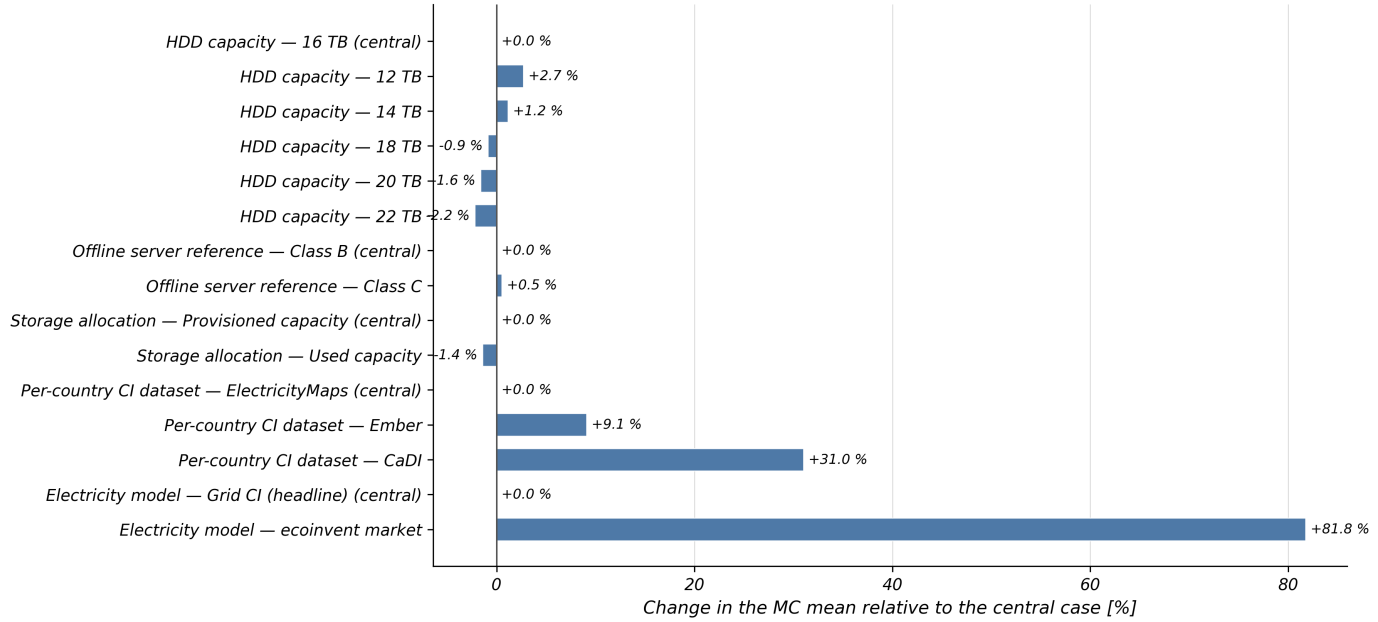


Figure 3.12: Effect of modelling variants on the Monte Carlo mean of the climate-change result in 2025

The choice of electricity data dominates the other variants. Replacing Electricity Maps carbon intensities with ecoinvent electricity mixes increases the mean result by 81.8 %. The differences produced by Ember, CaDI and ecoinvent do not reflect a physical change in the system under study; they arise from differences in reference year, modelling method and system boundary among the sources.

Hardware-related choices have a smaller effect on the total result. Allocating storage to the capacity actually occupied reduces the mean by 1.4 %, whereas using class C rather than class B to represent grid servers increases it by only 0.5 %. The latter result indicates that the choice of reference class server for offline embodied has little influence on the total footprint.

3.6 Comparative analysis: online versus offline computing

3.6.1 Shares and carbon intensity of computing

The comparison between electricity consumption and climate-change impact reveals an imbalance between the two systems, illustrated in the appendix (figure 20). The online branch accounts for 56 % of electricity use but 31 % of the climate-change impact; the corresponding offline shares are 44 % and 69 %. This difference results mainly from the higher carbon intensity of the electricity consumed by grid sites.

At the central parameter values, the impact of the complete online system per unit of provisioned HLT2 CPU capacity is 31.7 g CO_2eq per kHS23·h. The corresponding offline result is 156.5 g per kHS23·h of pledged CPU capacity, giving a ratio of 4.9 between the two reported indicators (figure 3.13). The whiskers propagate the uncertainty of the climate-impact numerator; they do not include uncertainty in the HS23 capacity denominators.

Figure 3.13 compares the climate-change impact per unit of provisioned computing capacity, expressed in g CO_2eq/(kHS23·h). It provides an average comparison of Monte Carlo production in the online and offline systems. Event Builder nodes are excluded from the online denominator because they include GPUs, which cannot currently run these Monte Carlo simulations.

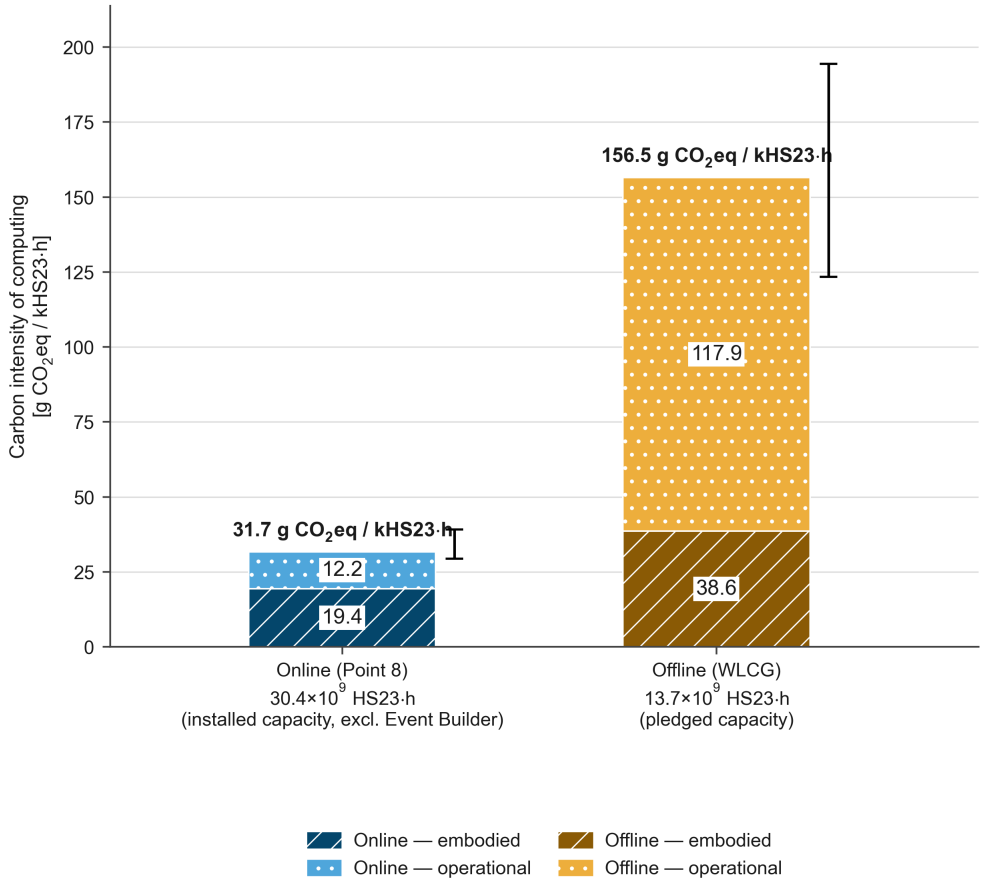


Figure 3.13: Climate-change impact per unit of provisioned CPU capacity — Electricity Maps

Their embodied and operational impacts remain in the numerator because the metered electricity cannot be separated by farm. Excluding their nominal CPU capacity reduces the denominator by 4.6 % and raises the reported intensity by approximately 4.8 % relative to an all-CPU-capacity denominator. The indicator therefore represents the impact of the complete online system per unit of provisioned HLT2 CPU capacity, rather than a pure measure of computing efficiency.

In 2025, the offline sites delivered approximately 19 % more computing work than their pledged capacity. The deterministic result per work actually delivered is 131.5 g CO_2eq per kHS23·h; the pledged-capacity value remains the basis of the comparison in figure 3.13.

3.6.2 Impact per unit of recorded luminosity

At the central parameter values, processing LHCb data in 2025 corresponds to an impact of 262.7 tonnes of CO_2eq per fb^{-1} based on recorded integrated luminosity (table 29). This result is based on the 11.8 fb^{-1} recorded during the year. This indicator follows the method defined in section 2.1 and therefore inherits its limitations and assumptions, including the steady-state assumption.

4. Discussion

4.1 Interpretation of the results

For 2025, the deterministic climate-change impact at the central parameter values is 3,100 t CO_2eq, including 2,137 t for the offline branch and 963 t for the online branch. Propagating the assigned uncertainties gives a Monte Carlo mean of 3,182 t and a P5–P95 interval of 2,644 t to 3,729 t. These values must be interpreted within the stated boundary.

The two branches have distinct profiles. Electricity represents about three quarters of the deterministic offline climate-change impact, whereas embodied impacts represent 61 % of the online total (figures 12 and 13 and table 28). This difference is primarily explained by electricity carbon intensity. The online data centre's PUE of 1.05 also limits its electricity use. When a data centre is supplied with low-carbon electricity, lower operational emissions increase the relative share of manufacturing, as observed for other computing facilities [56].

The reported impact per HS23·hour is approximately five times higher for the offline branch, but this ratio should not be interpreted as a direct comparison of hardware efficiency. The online denominator is installed HLT2 CPU capacity, the offline denominator is a contractual pledge, and the complete impact of the GPU-based Event Builder remains in the online numerator. The difference comes mainly from location and electricity supply, while grid averages also conceal large differences among sites.

Part of the HLT2 capacity is used for Monte Carlo simulation that would otherwise belong to the offline workload. When data-taking constraints leave HLT2 resources available, running these jobs at Point 8 makes use of already installed equipment and of electricity with a lower average carbon intensity than that of the grid. This can reduce idle capacity and displace work from offline sites.

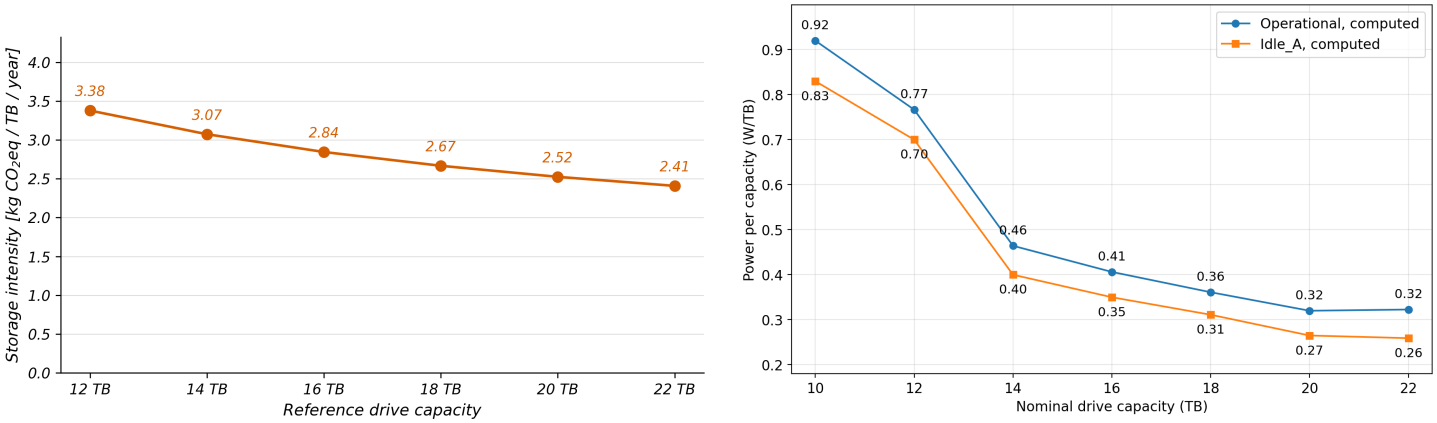
This distribution leads to different priorities. For the online branch, equipment lifetime and renewal decisions carry substantial weight. For the offline branch, the main levers first concern electricity consumption, equipment efficiency and the location of computing. This interpretation remains incomplete: the offline embodied estimate excludes both tape storage, which holds most of the data volume, and site infrastructure. It should therefore be treated as a lower bound.

The large manufacturing share of the online result also makes procurement relevant. The electricity supply and production conditions of semiconductor and equipment manufacturing can influence embodied impacts [19]. In the present model, generic processes and unknown manufacturing locations prevent a quantified comparison between countries or suppliers. Procurement decisions should consider manufacturing locations, since regional electricity mixes and production conditions can materially affect embodied emissions. contractual requirements may restrict supplier choice, requesting configuration-specific life-cycle data would enable more meaningful comparisons between suppliers.

The multi-impact profile calculated with ecoinvent also attributes the larger contribution to the offline branch in fourteen of the sixteen categories. Energy resources and ionising radiation are exceptions because of the contribution of French electricity to the online branch. Its low carbon intensity is accompanied by a higher contribution to the ionising-radiation indicator. The water-

use indicator includes upstream water consumption from electricity generation, other supply-chain processes and water consumed directly by the online data centre. Direct water consumption at offline sites is not covered.

The embodied impact of a drive per terabyte decreases as its capacity increases (figure 4.1a). Assessing whether this also reduces the total climate-change impact requires accounting for electricity use. The specific power of WD Gold drives decreases from 0.92 to 0.32 W/TB between 10 and 22 TB, a reduction of 65 % (figure 4.1b). At equal provisioned capacity, higher-capacity drives therefore reduce both the modelled embodied impact and electricity demand per terabyte.



(a) Embodied impact per terabyte and year (ResilioDB)

(b) Electrical power per terabyte (WD Gold datasheets)

Figure 4.1: Effect of individual drive capacity at constant total provisioned capacity: (a) embodied impact per terabyte and year; (b) electrical power per terabyte of WD Gold drives under operating and idle conditions

These differences demonstrate the value of a multi-impact interpretation. A climate-oriented measure may shift burdens towards water use, resources or toxicity. The scenario based on 22 TB rather than 16 TB disks reduces the modelled offline embodied impact by 12.1 % and the climate-change impact of the complete system by 2.2 % (tables 12 and 23). Fewer disks are required and their power per terabyte is lower.

For the servers, newer equipment may deliver more computing per unit of energy, but its manufacture adds an immediate impact. The appropriate replacement time therefore depends on the efficiency gain, utilisation, electricity carbon intensity and the impact of the new equipment. For the online data centre supplied with low-carbon electricity, extending use can be preferable when the expected operational saving does not compensate for the manufacturing impact during a realistic lifespan of a server[56].

The replacement threshold can be illustrated with the class B model. Replacing a functioning server creates an immediate embodied climate-change impact of approximately 2.27 t CO_2eq (table 7). This additional impact can be compensated only by the operational emissions avoided by the new equipment. Assuming ten years of continuous operation, a constant electricity intensity of 31 g CO_2eq/kWh and a PUE of 1.05, the avoided impact is calculated as $\Delta P \times 10 \times 8760 \times 1.05 \times 0.031$, where ΔP is the average reduction in IT power expressed in kW. Equating this saving to 2,270 kg CO_2eq gives a required reduction of approximately 0.8 kW. The replacement would therefore need to provide the same computing service while reducing average IT power by this amount throughout the ten-year period.

Finally, low PUE and low-carbon electricity do not eliminate absolute impacts. Upgrade II is expected

to increase event rates and the amount of data to process [8]. Absolute reductions may therefore become more difficult even if the impact per unit of computing or scientific service decreases. Both absolute impacts and performance-based indicators should be monitored. Otherwise, lower impacts per job can be offset by growth in the number of jobs or stored data, a rebound effect. Sobriety in this context does not mean limiting scientific output, but avoiding unnecessary recomputation, right-sizing retained data and ensuring that efficiency gains are not treated as room for more compute. User information and estimation tools can make the environmental cost of computing visible and encourage users to assess whether a calculation is necessary [23, 24]. Reducing electricity consumption would also lower operating fees, although quantifying that benefit would require a separate analysis.

4.2 Robustness of the results

The main patterns persist across the tested variants: electricity dominates the offline branch, manufacturing is important in the online branch, and the geography of the WLCG influences the climate-change impact. Absolute values nevertheless remain sensitive to the electricity model. Replacing Electricity Maps intensities with ecoinvent electricity mixes raises the mean result by 81.8 %, illustrating the differences in reference year, modelling method and system boundary between the two sources.

Model representativeness explains 70.5 % of the variance of the main climate indicator. Uncertainty therefore arises primarily from the description of the actual fleet. Repeating existing measurements would reduce it little; characterising the servers, power, age and storage models used at each site would be more useful. The embodied intensity of the class C proxy is only 13.7 % above that of class B, and substituting it changes the complete-system mean by 0.5 %. This supports the conclusion that the class B choice has little influence on the total result. It is not an independent validation of the server impacts, because both configurations are generated by the same ResilioDB modelling framework and share part of its underlying assumptions. Although hardware choices have only a modest effect on the total footprint, procurement and replacement remain relevant mitigation levers, particularly because the dominant impacts may be difficult to reduce. Reductions should therefore be pursued wherever feasible.

4.3 Limitations

4.3.1 Boundary and missing data

Coverage differs between the online and offline branches. The online inventory includes servers, storage, network equipment and the main infrastructure. For the offline branch, only computing servers and disk storage are modelled. Magnetic tapes, site buildings and technical installations, inter-site transfers and direct water use remain outside the boundary. Network impacts are therefore only partly covered: online network equipment is included, but the full equipment and energy required for data transfers is not.

The offline disk model also applies no additional capacity for RAID or erasure coding because the redundancy configuration of each site is unknown. If the WLCG provisioned-capacity figures describe usable rather than raw physical capacity, additional drives and their electricity use are omitted.

The origin of several components is unknown. Generic manufacturing processes are consequently assigned, and these do not necessarily represent the place, technology or year of manufacture.

This limitation contributes to representativeness uncertainty. A data-provenance map would be useful if it clearly separated measured information, assumptions and generic processes.

4.3.2 Modelling assumptions and comparability

The offline server lifetime is also assumed to be ten years, with the same uncertainty range as the online fleet. This central value comes from the on-site cascading and maintenance practice, not from observations at WLCG sites. Applying it to the distributed grid may overestimate the lifetime of servers operated under different maintenance and renewal policies. If their actual service lives are shorter, the annual offline embodied impact is underestimated, even though the six-to-twelve-year distribution partly covers this uncertainty.

Offline energy use is estimated from average models applied to the sites, whereas the online branch relies more extensively on measurements and a local inventory. Annual national carbon-intensity factors also represent neither site-specific electricity contracts nor local generation such as photovoltaic installations.

The ecoinvent electricity mixes also have limited temporal representativeness for 2025 operational conditions. Reconstructing the French mix with more recent data would not have been sufficient on its own, because the same reconstruction would have been required for every country represented in the offline system. For embodied impacts, older upstream energy data may be more representative of equipment manufactured before 2022 than of current production conditions.

The measured online PUE and WUE are likewise specific to the weather and IT load of 2025. Warmer outside conditions reduce the temperature difference available for free-air cooling and may require more fan electricity, more adiabatic water, or both. Climate warming and the higher heat load expected with Upgrade II could therefore change future PUE, WUE and associated impacts [3, 8]. This work does not project that interaction.

The comparison of the online and offline intensity indicators [gram CO_2eq per kHS23·h] also uses different denominators: installed capacity for the former and pledged capacity for the latter.

4.3.3 Temporal interpretation of the luminosity indicator

The indicator per unit of recorded luminosity (section 3.6.2) covers the complete chain and follows the annual accounting convention for the offline branch. An offline year's workload also depends on simulation and reprocessing campaigns scheduled by the collaboration. At identical configuration and efficiency, a year with a large campaign may therefore have a higher impact per fb^{-1} than a quieter year. The indicator measures the impact of processing performed during the year, not the final impact attributable only to data recorded in that same year.

4.4 Contributions of this work

This work brings life cycle assessment, data-centre operation and high-energy-physics computing needs into a common framework. It covers two levels of data processing, including a globally distributed infrastructure, together with embodied and operational impacts, sixteen environmental categories and uncertainty propagation. Reporting the results for a defined LHCb data-processing service, rather than for electricity or equipment alone, follows the service-based logic proposed for environmental assessments of IT services [57].

Multi-impact assessments at this scale remain uncommon and rely on dispersed information. This work gathers that information, states the required assumptions and proposes a method shared

across the three fields involved.

It provides an internal reference, identifies the dominant contributions and specifies the data that limit result quality. The method can be updated as site inventories, power profiles and storage data become available. Its practical contribution is a common framework for directing data collection and assessing reduction measures.

4.5 Possible improvements

Temporal scheduling of Monte Carlo simulations was examined using quarter-hourly French carbon intensities from 2019 to 2023 [58]. Recurrent variations are short and small, while HLT2 is often highly utilised. Under the conditions studied, moving jobs to lower-carbon periods does not appear to be a priority, although this option will be discussed internally. A marginal-electricity analysis could change this interpretation, but it would answer a consequential question different from the attributional approach used here.

Optimising processor frequency around the maximum HS23/W is another possible measure [32]. Processor power-saving states could complement it during idle or low-load periods. Both options should be tested on representative LHCb workloads and account for execution time, throughput and latency: lower instantaneous power does not necessarily reduce total energy if the calculation takes longer, and the benefit may be limited when utilisation is already high.

The work of Giemza, Padee, and Niewiadomska-Szynkiewicz [37] estimates the impacts of individual LHCb jobs and computing sites using data on processor performance and power consumption. This detailed approach complements the method developed in this thesis.

The most useful extension for the offline branch would be to collect, by site, the processor models that execute LHCb jobs (or at least, an inventory of their fleet) and the local PUE. These data would replace the uniform 5.5 W per logical core and PUE of 1.45.

Rainwater recovery could reduce mains-water withdrawal for cooling. Its relevance would depend on rainfall, storage volume, water-quality requirements, pumping energy and the impacts of the added infrastructure. A feasibility assessment is therefore required before it can become a recommendation.

Finally, a social or economic assessment could complement the environmental evaluation of this work.

5. Conclusions

This work assessed the environmental impacts of the LHCb data-processing chain over one reference year under RUN operating conditions. Within the stated system boundary, the assessment covers the online data centre and the offline resources of the WLCG, and combines embodied and operational impacts with a multi-impact and uncertainty analysis. It thereby provides a common basis for identifying the main environmental hotspots of the computing service delivered to the experiment.

For 2025, the deterministic climate-change impact at the central parameter values is estimated at 3,100 t CO_2eq, including 2,137 t for the offline branch and 963 t for the online branch. The Monte Carlo mean is 3,182 t, with a P5–P95 interval of 2,644 t to 3,729 t. The two branches have distinct profiles: electricity dominates the offline footprint, whereas hardware manufacture represents most of the online footprint. Disk storage is also a major contributor to offline embodied impacts. In the ecoinvent-based multi-impact profile, the offline branch contributes more than the online branch in fourteen of the sixteen impact categories. Energy resources and ionising radiation are the two exceptions, mainly because of the contribution of French electricity to the online branch.

The hierarchy of the main hotspots remains stable across the tested configurations, but the absolute result is sensitive to the electricity model. Model representativeness accounts for 70.5 % of the variance of the main climate indicator, showing that better knowledge of the actual equipment fleet would improve the assessment more than repeating already available measurements. The online–offline comparison must nevertheless be interpreted in light of the different levels of detail of the two inventories. The offline branch relies on top-down proxies and excludes tape storage, site infrastructure and direct water use; its embodied-impact estimate should therefore be regarded as a lower bound rather than as strictly comparable to the more detailed online inventory.

The priorities consequently differ between the two branches. The online data centre already benefits from low-carbon French electricity and a PUE of 1.05, which increases the relative importance of hardware manufacture. Purchases should therefore be limited to justified needs, and equipment lifetime should be extended when the operational savings from replacement do not compensate for the impact of manufacturing new equipment. When data-taking conditions allow, using the installed HLT2 capacity for Monte Carlo production also avoids leaving available resources idle and can displace work otherwise performed at offline sites. For the offline branch, the first priority is to collect site-specific information on processor models, power consumption, PUE and storage equipment, and then use it to target efficiency measures where the reduction potential is greatest. These actions should be accompanied by user awareness and by monitoring both absolute impacts and performance-based indicators. A multi-impact perspective must be retained so that reductions in climate-change impact do not create burdens in other environmental categories.

The online total is therefore not attributable solely to EB, HLT1 and HLT2 data taking, since the farm also runs Monte Carlo production during a substantial part of the year.

By assembling dispersed data and formalising the modelling choices, this work provides LHCb with a reproducible baseline and an updatable framework for future assessments. This is particularly relevant for Upgrade II, whose higher event rates and data volumes are expected to increase computing demand. Environmental performance will therefore need to be monitored alongside throughput and reliability, with attention to both efficiency gains and absolute impacts.

6. AI in this work

The AI tools used in this work were ChatGPT and Claude. AI was used occasionally to clarify concepts in particle physics and computing and to support literature searches. The author selected and consulted every cited source, either in full or, at least, in the sections required for its use. During method development, AI helped structure the reasoning on the functional unit and Monte Carlo simulations. It also assisted with code used to organise data and implement calculations. AI was not used to collect primary data. The author chose the assumptions, reviewed the code and verified the calculations. AI assisted with code used to prepare tables, graphs and figures. The interpretation of the results, the limitations and the recommendations were almost entirely developed by the author. AI was used to improve the structure, for proofreading and language. The author reviewed the resulting text and remains responsible for the content of the thesis.

References

[1] Millennium Ecosystem Assessment. *Ecosystems and Human Well-being: Synthesis*. Washington, DC: Island Press, 2005. ISBN: 978-1-59726-040-4. URL: https://www.millenniumassessment.org/documents/document.356.aspx.pdf.

[2] Katherine Richardson et al. "Earth beyond six of nine planetary boundaries". In: *Science Advances* 9.37 (2023), eadh2458. DOI: 10.1126/sciadv.adh2458. URL: https://doi.org/10.1126/sciadv.adh2458.

[3] IPCC. *Climate Change 2023: Synthesis Report*. Geneva, Switzerland: Intergovernmental Panel on Climate Change, 2023. DOI: 10.59327/IPCC/AR6-9789291691647. URL: https://www.ipcc.ch/report/ar6/syr/.

[4] International Energy Agency. *Energy and AI*. Licence: CC BY 4.0. Paris, France: International Energy Agency, 2025. URL: https://www.iea.org/reports/energy-and-ai.

[5] Kenneth Bloom et al. *Climate impacts of particle physics*. White paper submitted to the Snowmass 2021 community planning exercise. 2022. DOI: 10.48550/arXiv.2203.12389. arXiv: 2203.12389 [hep-ex]. URL: https://arxiv.org/abs/2203.12389.

[6] European Organization for Nuclear Research (CERN). *Vol. 4 (2025): CERN Environment Report — Rapport sur l'environnement 2023–2024*. CERN-Environment-2025-004. CERN, 2025. DOI: 10.25325/CERN-Environment-2025-004. URL: https://doi.org/10.25325/CERN-Environment-2025-004.

[7] European Organization for Nuclear Research (CERN). *CERN Energy Overview*. Web application providing energy consumption and electricity sourcing data for CERN facilities. 2026. URL: https://energy.app.cern.ch/overview (visited on 06/08/2026).

[8] LHCb Collaboration. *LHCb Upgrade II Scoping Document*. Technical Design Report CERN-LHCC-2024-010; LHCB-TDR-026. Geneva, Switzerland: CERN, 2024. DOI: 10.17181/CERN.2RXP.HDK0. URL: https://cds.cern.ch/record/2903094.

[9] Lukas Calefice et al. "Effect of the high-level trigger for detecting long-lived particles at LHCb". In: *Frontiers in Big Data* 5 (2022), p. 1008737. ISSN: 2624-909X. DOI: 10.3389/fdata.2022.1008737. URL: https://doi.org/10.3389/fdata.2022.1008737.

[10] Worldwide LHC Computing Grid Collaboration. *Memorandum of Understanding for Collaboration in the Deployment and Exploitation of the Worldwide LHC Computing Grid*. CERN document CERN-C-RRB-2005-01, as amended. Defines the Tier structure and the pledge mechanism of the WLCG. 2005. URL: https://wlcg.web.cern.ch/collaboration/mou.

[11] Statista. *Data center average annual power usage effectiveness (PUE) worldwide from 2007 to 2025*. Statistic based on global data center average annual PUE values. 2025. URL: https://www.statista.com/statistics/1229367/data-center-average-annual-pue-worldwide/ (visited on 07/20/2026).

[12] ODISSEE Project. *ODISSEE Project: Using AI to cope with data deluge from SKAO and CERN HL-LHC*. Official project website. 2026. URL: https://odissee-project.eu/ (visited on 07/17/2026).

[13] Niko Neufeld et al. *Sustainable Computing Solutions: A Case-Study of the LHCb Data-Centre*. Presentation at ICHEP 2024, Prague. Slides available on CERN Indico. 2024. URL: https:

//indico.cern.ch/event/1291157/contributions/5900404/attachments/2898960/5083248/sustainable_cooling_solutions_LHCb_N_Neufeld_ICHEP204.pdf.

[14] *Environmental management — Life cycle assessment — Principles and framework*. 2nd ed. ISO 14040:2006. Published 2006-07. International Organization for Standardization. July 2006. URL: https://www.iso.org/standard/37456.html.

[15] *Environmental management — Life cycle assessment — Requirements and guidelines*. 1st ed. ISO 14044:2006. International Organization for Standardization. July 2006. URL: https://www.iso.org/standard/38498.html.

[16] *Information technology — Data centres key performance indicators — Part 2: Power usage effectiveness (PUE)*. 2nd ed. ISO/IEC 30134-2:2026. International Organization for Standardization and International Electrotechnical Commission. Jan. 2026. URL: https://www.iso.org/standard/30134-2.

[17] *Information technology — Data centres key performance indicators — Part 9: Water usage effectiveness (WUE)*. ISO/IEC 30134-9:2022. International Organization for Standardization and International Electrotechnical Commission. Mar. 2022. URL: https://www.iso.org/fr/standard/77692.html.

[18] International Telecommunication Union. *Methodology for Environmental Life Cycle Assessments of Information and Communication Technology Goods, Networks and Services*. ITU-T L.1410. Recommendation ITU-T L.1410, in force version 11/2024. International Telecommunication Union. Nov. 2024. URL: https://www.itu.int/rec/T-REC-L.1410.

[19] Thibault Pirson et al. "The Environmental Footprint of IC Production: Review, Analysis, and Lessons From Historical Trends". In: *IEEE Transactions on Semiconductor Manufacturing* 36.1 (2023), pp. 56–67. DOI: 10.1109/TSM.2022.3228311. URL: https://doi.org/10.1109/TSM.2022.3228311.

[20] ATLAS Collaboration. "The Environmental Impact, Carbon Emissions and Sustainability of Computing in the ATLAS Experiment". In: *The European Physical Journal C* 85 (2025), p. 1397. DOI: 10.1140/epjc/s10052-025-14976-3. URL: https://doi.org/10.1140/epjc/s10052-025-14976-3.

[21] Natalia Diana Szczepanek and Domenico Giordano. *Power and Efficiency Monitoring for WLCG Sites*. Presentation at the HEPiX Spring 2026 Workshop, Lisbon. The related dataset of 212 processor models used in this work was supplied to the author and is not published separately. Apr. 2026. URL: https://indico.cern.ch/event/1598655/contributions/7011092/ (visited on 08/10/2026).

[22] David Britton, Simone Campana, and Bernd Panzer-Steindel. "A holistic study of the WLCG energy needs for the LHC scientific program". In: *EPJ Web of Conferences* 295 (2024), p. 04001. DOI: 10.1051/epjconf/202429504001. URL: https://doi.org/10.1051/epjconf/202429504001.

[23] Loïc Lannelongue, Jason Grealey, and Michael Inouye. "Green Algorithms: Quantifying the Carbon Footprint of Computation". In: *Advanced Science* 8.12 (2021), p. 2100707. DOI: 10.1002/advs.202100707. URL: https://doi.org/10.1002/advs.202100707.

[24] Loïc Lannelongue et al. "GREENER principles for environmentally sustainable computational science". In: *Nature Computational Science* 3.6 (2023), pp. 514–521. DOI: 10.1038/s43588-023-00461-y. URL: https://doi.org/10.1038/s43588-023-00461-y.

[25] GreenDelta. *openLCA*. Version 2.6.2. Life cycle assessment software; this release corrects inclusion of linked upstream results in Monte Carlo simulation. June 10, 2026. URL: https://www.openlca.org/download/ (visited on 08/09/2026).

[26] ecoinvent Association. *ecoinvent database*. Version 3.12. 2025. URL: https://support.ecoinvent.org/ecoinvent-version-3.12 (visited on 08/09/2026).

[27] Resilio. *ResilioDB*. Version 2026.7. Generic portal for the specialised life cycle inventory and impact-modelling database for information technology hardware; an authenticated account is required to inspect database content. 2026. URL: https://db.resilio.tech/ (visited on 08/10/2026).

[28] Electricity Maps. *Yearly Carbon-Intensity Data for 2025*. Version 2025 yearly dataset redistributed in @tgwf/co2 0.18.0. Annual location-based consumption carbon intensities used in the operational climate-change calculations. 2025. URL: https://www.electricitymaps.com/data/methodology (visited on 08/09/2026).

[29] The Green Web Foundation. *CO2.js*. Version 0.18.0. Version used to redistribute the Electricity Maps 2025 yearly dataset. 2026. URL: https://www.npmjs.com/package/@tgwf/co2/v/0.18.0 (visited on 08/09/2026).

[30] Michael Z. Hauschild, Ralph K. Rosenbaum, and Stig Irving Olsen. *Life Cycle Assessment: Theory and Practice*. Cham, Switzerland: Springer International Publishing, 2018. ISBN: 978-3-319-56475-3. DOI: 10.1007/978-3-319-56475-3. URL: https://link.springer.com/book/10.1007/978-3-319-56475-3.

[31] Open Compute Project Foundation. *Data Center Facility Sustainability Metrics*. Tech. rep. Version 1.0. OCP Data Center Facilities Project. Open Compute Project Foundation, Apr. 2023. URL: https://www.opencompute.org/documents/data-center-facility-sustainability-metrics-white-paper-v1-0-pdf.

[32] Natalia Szczepanek et al. *HEP Benchmark Suite: Enhancing Efficiency and Sustainability in Worldwide LHC Computing Infrastructures*. 2024. DOI: 10.48550/arXiv.2408.12445. arXiv: 2408.12445 [cs.DC]. URL: https://arxiv.org/abs/2408.12445.

[33] European Commission. *Commission Recommendation (EU) 2021/2279 on the use of the Environmental Footprint methods to measure and communicate the life cycle environmental performance of products and organisations*. Official Journal of the European Union, L 471, 15.12.2021, pp. 1–396. Product Environmental Footprint (PEF) method, version 3.1. 2021. URL: https://eur-lex.europa.eu/eli/reco/2021/2279/oj.

[34] Susanna Andreasi Bassi et al. *Updated Characterisation and Normalisation Factors for the Environmental Footprint 3.1 Method*. JRC Technical Report EUR 31414 EN. JRC130796. Luxembourg: European Commission, Joint Research Centre, 2023. DOI: 10.2760/798894. URL: https://publications.jrc.ec.europa.eu/repository/handle/JRC130796 (visited on 08/07/2026).

[35] Mattia Damiani, Nicola Ferrara, and Fulvio Ardente. *Understanding Product Environmental Footprint and Organisation Environmental Footprint Methods*. JRC Technical Report EUR 31236 EN. JRC129907. Luxembourg: European Commission, Joint Research Centre, 2022. DOI: 10.2760/11564. URL: https://publications.jrc.ec.europa.eu/repository/handle/JRC129907.

[36] EGI Foundation. *Accounting Portal*. Documentation of the EGI accounting data, including CPU and wall-clock time and work metrics. 2026. URL: https://docs.egi.eu/internal/accounting/portal/ (visited on 08/17/2026).

[37] Henryk Giemza, Adam Padee, and Ewa Niewiadomska-Szynkiewicz. *Estimating the Carbon Footprint of Computations in the LHCb Grid*. Presentation at the 28th Conference on Computing in High Energy and Nuclear Physics (CHEP 2026). Slide 6 reports an average CPU utilisation of 96 percent for LHCb jobs. May 2026. URL: https://indico.cern.ch/event/1471803/contributions/6966469/ (visited on 08/10/2026).

[38] GIGABYTE Technology. *G482-Z51 (Rev. A00) GPU Server*. Technical specifications consulted for server configuration and hardware inventory. 2025. URL: https://www.gigabyte.com/Enterprise/GPU-Server/G482-Z51-rev-A00 (visited on 06/03/2026).

[39] Super Micro Computer, Inc. *CloudDC SuperServer SYS-621C-TN12R*. Technical specifications consulted for server configuration and hardware inventory. 2025. URL: https://www.supermicro.com/en/products/system/clouddc/2u/sys-621c-tn12r (visited on 06/03/2026).

[40] Roel Aaij et al. "Allen: A High-Level Trigger on GPUs for LHCb". In: *Computing and Software for Big Science* 4.1 (2020), p. 7. DOI: 10.1007/s41781-020-00039-7. URL: https://doi.org/10.1007/s41781-020-00039-7.

[41] Niko Neufeld and Heinrich Schindler. *The LHCb Online Data Centre at Point 8: Design, Operation and Cooling*. LHCb Online internal documentation. Internal LHCb Online documentation; add an EDMS record, internal note or Indico presentation before citing this entry. 2025.

[42] European Organization for Nuclear Research (CERN). *Data Center Cumulative PUE*. CERN Grafana dashboard for cumulative Power Usage Effectiveness (PUE) monitoring. 2026. URL: https://lbgrafana.cern.ch/d/VFRTrIZVz/data-center-cumulative-pue?orgId=1 (visited on 07/20/2026).

[43] European Organization for Nuclear Research (CERN). *Building Work for CERN's New Data Centre in Prévessin Begins*. Reports a design PUE target of approximately 1.1 for Prévessin and a PUE of approximately 1.5 for Meyrin. Apr. 2022. URL: https://home.cern/news/news/computing/building-work-cerns-new-data-centre-prevessin-begins (visited on 08/09/2026).

[44] Mary Sotos. *GHG Protocol Scope 2 Guidance: An Amendment to the GHG Protocol Corporate Standard*. Guidance document. Published by the Greenhouse Gas Protocol. World Resources Institute, 2015. URL: https://ghgprotocol.org/sites/default/files/2023-03/Scope%202%20Guidance.pdf.

[45] Electricity Maps. *Why Marginal Emission Factors Are Unsuitable for Scope 2 Accounting*. Online publication discussing the limitations of marginal emission factors for Scope 2 carbon accounting. 2024. URL: https://www.electricitymaps.com/resources/publications/why-marginal-emission-factors-are-unsuitable-for-scope-2-accounting (visited on 07/20/2026).

[46] Luiz André Barroso and Urs Hölzle. "The Case for Energy-Proportional Computing". In: *Computer* 40.12 (2007), pp. 33–37. DOI: 10.1109/MC.2007.443. URL: https://doi.org/10.1109/MC.2007.443.

[47] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz André Barroso. "Power Provisioning for a Warehouse-Sized Computer". In: *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA '07)*. New York, NY: ACM, 2007, pp. 13–23. DOI: 10.1145/1250662.1250665. URL: https://doi.org/10.1145/1250662.1250665.

[48] European Organization for Nuclear Research (CERN). *CERN Signs Long-Term Solar Power Agreements*. States that CERN electricity is, under normal circumstances, procured entirely from France. Nov. 2024. URL: https://home.cern/cern-signs-long-term-solar-power-agreements/ (visited on 08/09/2026).

[49] Seagate Technology. *Exos X16 SATA Product Manual*. Seagate Technology. 2019. URL: https://www.seagate.com/files/www-content/product-content/enterprise-hdd-fam/exos-x-16/en-us/docs/100855249a.pdf (visited on 08/14/2026).

[50] Seagate Technology. *Exos X18 SATA Product Manual*. Seagate Technology. 2020. URL: https://www.seagate.com/files/www-content/product-content/enterprise-hdd-fam/exos-x18/_shared/en-us/docs/100865854a.pdf (visited on 08/14/2026).

[51] Europe Economics. *Impact of Growth of Data Centres on Energy Consumption*. Tech. rep. Table 3.3 reports a central fixed-network core-and-access intensity of 0.0016 kWh/GB. UK Department for Energy Security and Net Zero, 2025. URL: https://assets.publishing.service.gov.uk/media/689d9dc487bf475940723f6c/impact-of-growth-of-data-centres.pdf (visited on 08/10/2026).

[52] LHCb Collaboration. *LHCb Transfer Throughput, 2025*. Internal Grafana export. Hourly site-level throughput series exported on 16 July 2026; the dataset is not publicly released. 2025.

[53] Andreas Ciroth et al. "Empirically based uncertainty factors for the pedigree matrix in ecoinvent". In: *The International Journal of Life Cycle Assessment* 21.9 (2016), pp. 1338–1348. DOI: 10.1007/s11367-013-0670-5. URL: https://doi.org/10.1007/s11367-013-0670-5.

[54] Resilio. *ResilioDB — Data Quality Scoring Methodology*. Version 2026.7. Authenticated documentation for the database version used in this work; readers without a ResilioDB account cannot inspect the documentation directly; scores range from 0 to 5. 2026. URL: https://db.resilio.tech/docs (visited on 08/10/2026).

[55] Shannon M. Lloyd and Robert Ries. "Characterizing, Propagating, and Analyzing Uncertainty in Life-Cycle Assessment: A Survey of Quantitative Approaches". In: *Journal of Industrial Ecology* 11.1 (2007), pp. 161–179. DOI: 10.1162/jiec.2007.1136. URL: https://doi.org/10.1162/jiec.2007.1136.

[56] Mattias Wadenstein and Wim Vanderbauwhede. "Life cycle analysis for emissions of scientific computing centres". In: *The European Physical Journal C* 85 (2025), p. 913. DOI: 10.1140/epjc/s10052-025-14650-8. arXiv: 2506.14365. URL: https://doi.org/10.1140/epjc/s10052-025-14650-8.

[57] Julia Paolini et al. *MITSI – Methodology for IT Services Environmental Impact Assessment*. Technical report. Version 1.0. École Polytechnique Fédérale de Lausanne (EPFL), Mar. 2026. DOI: 10.5075/epfl.20.500.14299/261022. URL: https://infoscience.epfl.ch/entities/publication/39b2f2d9-aae1-4d12-88f2-dc15e633d8b1.

[58] RTE – Réseau de Transport d'Électricité. *Download éCO2mix Data*. Accessed: 2026-06-03. 2026. URL: https://www.rte-france.com/en/data-publications/eco2mix/download-indicators.

[59] CERN. *EDMS Document No. 100243*. https://edms.cern.ch/document/100243/2.2. CERN Engineering Data Management System (EDMS). 2004. URL: https://edms.cern.ch/document/100243/2.2.

Appendices

Methodological supplements

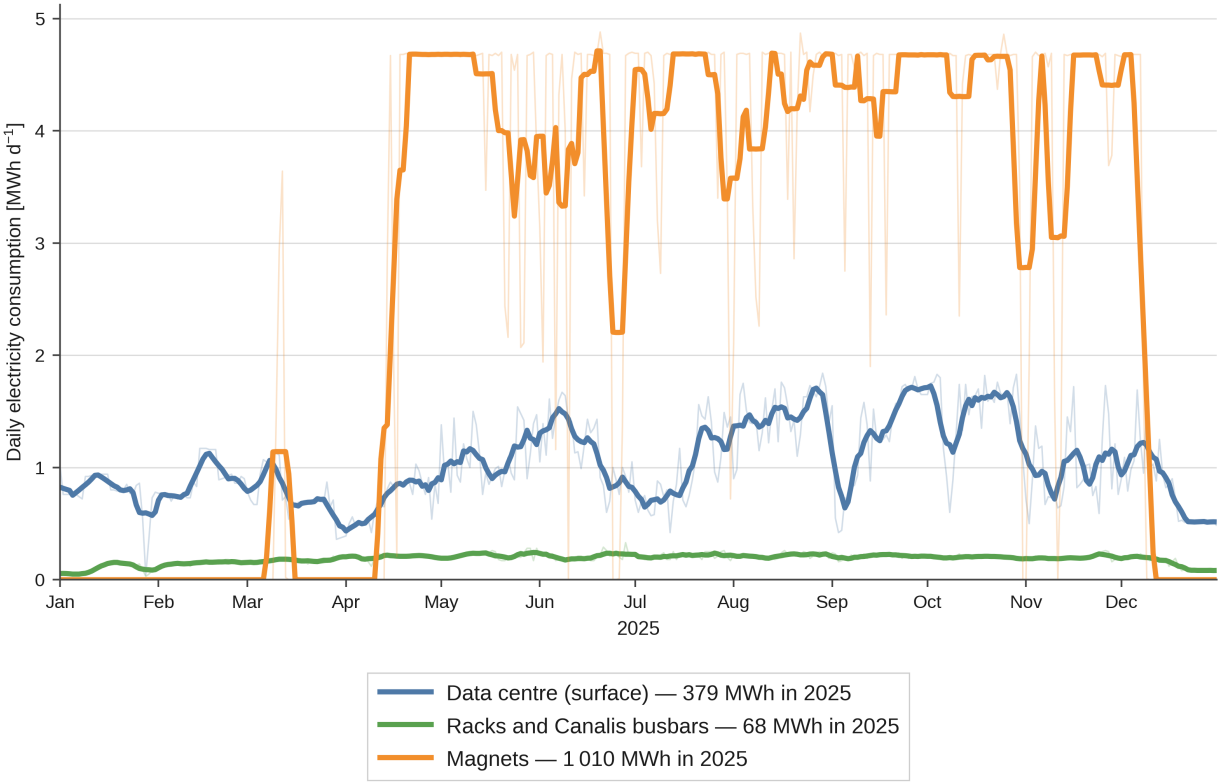


Figure 1: Daily average LHCb electrical power in 2025

Table 1: Assessment data, models and outputs

Branch	Nature	Foreground data	Background models and factors	Model output
Online	Embodied	Installed equipment inventory: servers, storage, network, cooling and electrical equipment, containers, foundations, racks and cables; configurations, quantities, material compositions and assigned lifetimes.	ResilioDB 2026.7 for servers, HDDs and routers; openLCA 2.6.2 with ecoinvent 3.12 APOS for non-IT equipment; component-specific or default end-of-life scenarios.	Cradle-to-grave EF 3.1 impacts annualised by lifespan, equipment family and life-cycle stage.
Online	Operational	Point 8 measurements for 2025: total, IT and air-handling-unit electricity; daily water withdrawals; French electricity supply.	Electricity Maps 2025 France for the headline climate result; ecoinvent 3.12 French electricity mix and water processes for the sixteen-category profile.	Annual electricity and water indicators, PUE and WUE, and operational EF 3.1 impacts.
Offline	Embodied	WLCG wall-clock occupation in logical core-hours; class B reference server (class C sensitivity); monthly provisioned disk capacity; 16 TB reference HDD (12–22 TB variants); equipment lifetimes.	ResilioDB server and HDD models with default end-of-life treatment.	Annual embodied impacts allocated to compute occupation and provisioned storage; lower-bound result because tape and site infrastructure are excluded.
Offline	Operational	WLCG wall-clock occupation by site; $P_{\text{core}} = 5.5$ W; PUE of 1.45; monthly disk capacity by site; disk specific power of 0.41 W/TB; country electricity attribution.	Electricity Maps 2025 country intensities for climate change; ecoinvent 3.12 country mixes for the sixteen-category profile.	Compute and disk electricity and operational impacts by site and country; water, tape and inter-site transfers are excluded.

Cables

The cables of the online data centre transfer information as electrical signals between servers, routers and switches or as optical signals from the detector to the surface data centre.

The inventory is based on a file listing all cables recorded in the cavern and the online data centre. It contains more than 23,000 cables, each with its own identifier, of which slightly fewer than 4,000 are attributed to the online system. Only the online cables are reported in this work.

The online cables were identified and selected on the basis of their respective identifiers, each cable being labelled according to the CERN convention in force [59].

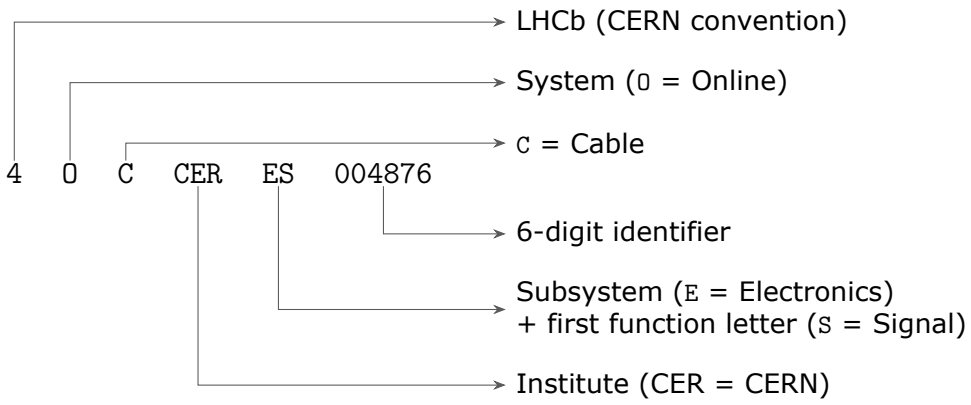


Figure 2: CERN cable-numbering convention

The number of cables and the diversity of their compositions and dimensions do not allow a fine, individual analysis of their composition. In order to build the bill of materials of such a large population, the cables were therefore grouped by function. All cables, whatever their function, share a similar architecture: each cable is divided into a conducting part, a protective part (shielding) and an insulating part. Each function was then assigned an estimate of the cross-section surface shares of these three parts. The shielding was not considered further: by engineering judgement agreed with the on-site teams, its cross-sectional share was assumed negligible compared with the conducting surface.

The volumes of the cables were then evaluated from their diameter and length, information also available in the initial file. This enabled an aggregation and the computation of a global cable volume (and of its three parts) for each function.

Table 2: Cable material composition

Function	Material			Ratio [% of section surface]		
	Conducting	Shielding	Insulation	Conducting	Shielding	Insulation
Signal	Copper	Copper	Polyolefin	10	/	90
TFC	Copper	Copper	Polyolefin	5	/	95
ECS	Copper	Copper	Polyolefin	5	/	95

Modelling of the remaining online infrastructure

In addition to the cables and IT equipment, the foreground inventory of the online system includes the prefabricated modules, their foundations, the air-handling units (AHUs), water and reverse-osmosis pumps, inverters, transformers, racks and air filters. These components were modelled in

openLCA 2.6.2 using background processes from ecoinvent 3.12 under the APOS system model. The installed quantities and reference lifetimes are reported in tables 3 and 2.2. Cradle-to-grave impacts were calculated for each component and annualised over its assigned lifetime.

Two modelling approaches were used, depending on the available technical information. Where a sufficiently detailed bill of materials could be established, the masses of the constituent materials were entered as foreground flows and linked to the corresponding ecoinvent datasets. Where product-specific composition data were unavailable, the closest functionally equivalent ecoinvent dataset was used as a proxy. The same generic 22 kW water-pump dataset was used for both the water pumps and the reverse-osmosis pumps, while the inverters were represented by a generic 2.5 kW inverter dataset. Both proxies were scaled to the number of installed units. The air filters were represented by the ecoinvent dataset for an air filter used in a central air-handling unit, with a reference airflow of 600 m^3/h.

Table 3: Remaining online-infrastructure inventory

Component	**Model basis**	**Quantity**	**Life [yr]**
Containers	Steel, aluminium, polypropylene, polyurethane insulation, stainless steel, rubber, welding and zinc coating	7	50
Foundations	Reinforced and unreinforced concrete with reinforcing steel	1 system	50
Air-handling units	Steel and aluminium casing, copper and stainless-steel parts, mineral and glass wool, polymers and rubber	24	20
Air filters	Generic central-unit air-filter proxy rated at 600 m^3/h	192	0.5
Water and reverse-osmosis pumps	Ecoinvent 22 kW water-pump proxy	24 + 4	20
Inverters	Ecoinvent 2.5 kW inverter proxy	144	20
Transformers	Silicon steel, low-alloy steel, aluminium, epoxy resin, copper and brass	2	50
Racks	120 kg of low-alloy steel and sheet rolling per rack	132	50
Network equipment	Generic ResilioDB router/switch and InfiniBand-controller models	318 + 600	10

The foundations were modelled from the construction drawings of the online data centre.

The AHUs, transformers and racks were modelled from internal technical documentation. The available drawings, equipment specifications and mass data were converted into component-level bills of materials. For one AHU, the model contains approximately 2.62 t of unalloyed steel, 0.27 t of aluminium, 0.04 t of copper, 0.02 t of stainless steel, 0.18 t of stone wool, 0.11 t of glass wool and smaller quantities of epoxy resin, polycarbonate and synthetic rubber. One transformer contains approximately 3.18 t of grain-oriented silicon steel, 1.81 t of low-alloy steel, 0.90 t of aluminium, 1.55 t of epoxy resin, 0.13 t of copper and 0.04 t of brass. Each rack was represented by 120 kg of low-alloy steel together with the associated sheet-forming process.

The internal documents were accessed through the LHCb project area of CERN's Engineering and Equipment Data Management Service (EDMS).[1] The principal records available for this inventory in-

[1] LHCb project documentation in CERN EDMS.

clude the LHCb Data Centre infrastructure proposal (EDMS 1608894, revision 3.0), the Automation prefabricated-data-centre technical description D-150153B1, the STULZ CH2-S4 technical specification dated 15 November 2019, the TESAR transformer drawings under reference TRP-025-3150-0018-00 and the Panduit Net-Access S-Type rack specifications.

Public environmental product declarations (EPDs) for related HVAC, filtration and electrical equipment were used only as qualitative plausibility checks on material composition and orders of magnitude. The records consulted were the Midea split-type room air conditioner (S-P-06897), the Carrier FJ4 fan-coil and 38MURA heat-pump units (EPD-IES-0016210), the Ultramare KULING air filter (EPD-IES-0017325:001), the Toshiba 8174 transformer series (S-P-06464), the Toshiba 8574C transformer (S-P-09217), and the Shirdi Sai 30 kV A and 200 kV A transformers (EPD-IES-0020031:001 and EPD-IES-0020036:001). These products differ from the installed equipment in scale, technology and geography. They were consulted only for qualitative context and are not treated as directly comparable results.

The material breakdown of the seven prefabricated modules was estimated by combining the internal technical documentation with published information on the construction of maritime containers and comparable modular enclosures. For one module, the model contains approximately 27.21 t of low-alloy steel, 3.28 t of aluminium, 2.34 t of polypropylene, 4.49 t of rigid polyurethane foam, 0.10 t of chromium steel and 0.05 t of synthetic rubber. Arc welding and zinc coating were included as forming and surface-treatment processes.

The network equipment followed the IT-hardware modelling route rather than the openLCA infrastructure route. The 318 routers and switches and the 600 InfiniBand controllers were represented by the corresponding generic models in ResilioDB 2026.7 and annualised over the ten-year reference lifetime used for IT equipment. This treatment is included in the inventory and results reported in tables 3 and 7.

The treatment of transport, installation, replacement and end of life follows the modelling defined for the online embodied assessment in section 2.3.2. The main limitations of these models are the use of proxy datasets and the technological representativeness. These limitations are addressed through the data-quality and representativeness factors used in the uncertainty analysis.

Replacement scales for disks and servers

The following values place replacement decisions on the same GWP100 scale as the results. They describe embodied impacts only: operational energy savings, transport variations and failure-rate effects must be assessed separately before deciding whether replacement is beneficial.

The ResilioDB disk model is equivalent to approximately 0.455 t CO_2eq for one 16 TB HDD. Replacing 1,000 such drives therefore represents about 455 t CO_2eq. The reference storage fleet of 9,144 drives represents approximately 4,162 t CO_2eq over manufacture and end of life, or 416 t CO_2eq per year with the assumed ten-year life. Providing the same capacity with 22 TB drives requires about 6,650 units and represents roughly 3,524 t CO_2eq before annualisation.

For the two inventoried server configurations, the cradle-to-grave climate-change impact is approximately 2.24 t to 2.27 t CO_2eq per server. A replacement campaign of 200 comparable servers therefore represents about 447 t to 453 t CO_2eq.

Table 4: Monte Carlo parameters and distributions

Parameter	Scope	Distribution	Bounds	Group
Lifespan — IT	IT hardware: servers A/B/C, on-site storage, switches, InfiniBand, grid disks	Triangular	(6 ; 10 ; 12)	G4 — lifespans
Lifespan — long-lived	Containers, foundations, racks, cables, transformers	Triangular	(50 ; 50 ; 100)	G4 — lifespans
Lifespan — electromechanical	Air handling units, inverters, water and RO pumps	Triangular	(15 ; 20 ; 25)	G4 — lifespans
Lifespan — filters	Air filters	Deterministic	0.5	G4 — lifespans
Fleet size (multiplicative factor)	servers_C (central: 1225 units)	Triangular	(0.95 ; 1 ; 1.05)	G3 — input data
P_{core}	Grid compute electricity	Triangular	(2.9 ; 5.5 ; 7.9)	G2 — representativeness
PUE of grid sites	Grid compute and disk electricity	Triangular	(1.4 ; 1.45 ; 1.5)	G3 — input data
Disk specific power w	Grid disk electricity	Triangular	(0.3 ; 0.406 ; 0.6)	G2 — representativeness
Pedigree factor f_{data}	openLCA components with a global pedigree	Lognormal, median 1	GSD 1.05–1.49	G3 — input data
Pedigree factor f_{repr}	openLCA components with a global pedigree	Lognormal, median 1	GSD 1.05–2.00	G2 — representativeness
ResilioDB factor (manufacture source)	IT hardware, on-site and grid	Lognormal, median 1	GSD 1.05–1.09	G1 — dataset quality
ResilioDB factor (environmental assessment quality)	IT hardware, on-site and grid	Lognormal, median 1	GSD 1.05–1.20	G2 — representativeness
ResilioDB factor (input data quality)	IT hardware, on-site and grid	Lognormal, median 1	GSD 1.00–1.02	G3 — input data
Background exchange uncertainties	Three retained openLCA chains	Native (mostly lognormal)	As published	G1 — dataset quality

Provenance of the openLCA uncertainty chains

In April 2026, GreenDelta released corrected ecoinvent 3.12 database files for openLCA. The earlier files assigned only basic uncertainty to exchanges and omitted the pedigree contribution, which underestimated background dispersion. openLCA 2.6.2 also corrected a defect that prevented linked upstream results from being included in Monte Carlo simulations. All chains produced with the earlier database or software were discarded and recomputed.

Each retained chain records the ecoinvent database version, openLCA version, container-image identifier, uncertainty mode, requested and completed iterations, and timestamps. The assembly checks this provenance before calculation and stops if the chains do not reference the same database. Process uncertainty settings were restored from fresh snapshots after simulation and audited against the archived inventory.

Impact categories of the EF 3.1 method

Table 5: EF 3.1 impact categories [33–35]

Impact category	Impact-category indicator	Unit	Interest and interpretation
Climate change, total	Global warming potential over 100 years (GWP_{100})	kg CO_2 eq	Measures the contribution of greenhouse-gas emissions to climate change over a 100-year horizon. The unit converts the radiative effect of each gas into the mass of CO_2 that would produce an equivalent effect, which makes gases with different lifetimes and strengths comparable.
Ozone depletion	Ozone depletion potential (ODP)	kg CFC-11 eq	Tracks emissions that destroy stratospheric ozone and thereby increase harmful ultraviolet radiation at the Earth's surface. Results are expressed relative to the ozone-depleting effect of CFC-11.
Human toxicity, cancer	Comparative Toxic Unit for humans (CTU_h)	CTU_h	Estimates potential cancer cases associated with human intake of toxic substances emitted to air, water and soil. The comparative toxic unit combines environmental fate, exposure and cancer effect, allowing different substances to be added on a common potential-disease-case scale; it is not a measured clinical incidence.
Human toxicity, non-cancer	Comparative Toxic Unit for humans (CTU_h)	CTU_h	Estimates potential non-cancer disease cases associated with human intake of toxic substances emitted to air, water and soil. CTU_h places chemicals with different fate, exposure and effect profiles on one comparable potential-disease-case scale; direct occupational or consumer exposure is not automatically represented.
Particulate matter	Impact on human health	Disease incidence	Estimates additional disease cases caused by primary fine particles and by precursors that form particles in the atmosphere, such as SO_2, NO_x and NH_3. The unit reports a modelled number of cases rather than a mass of particles, so emissions with different health effects can be compared.
Ionising radiation, human health	Human exposure efficiency relative to U-235	kBq U-235 eq	Represents the potential human-health effect of releases of radionuclides. Activities of different radionuclides and exposure pathways are converted to the equivalent activity of U-235, expressed in kilobecquerels.

Continued on next page

Table 5 – continued from previous page

Impact category	Impact-category indicator	Unit	Interest and interpretation
Photochemical ozone formation, human health	Tropospheric ozone concentration increase	kg NMVOC eq	Quantifies the formation of harmful ground-level ozone (summer smog) from NO_x, carbon monoxide and volatile-organic-compound emissions in sunlight. The unit expresses their ozone-forming potential as an equivalent mass of non-methane volatile organic compounds (NMVOC).
Acidification	Accumulated exceedance (AE)	mol H^+ eq	Assesses the exceedance of ecosystem critical loads caused mainly by emissions of sulfur and nitrogen compounds. Expressing acidifying potential as moles of H^+ equivalent provides a common chemical basis while retaining the site-dependent sensitivity embedded in the characterisation factors.
Eutrophication, terrestrial	Accumulated exceedance (AE)	mol N eq	Measures nitrogen enrichment of terrestrial ecosystems beyond their critical loads, which can alter species composition and ecosystem functioning. The equivalent-mole-of-nitrogen unit compares the eutrophying effect of different airborne nitrogen compounds.
Eutrophication, freshwater	Fraction of nutrients reaching the freshwater end compartment (P)	kg P eq	Represents nutrient enrichment of rivers and lakes, where phosphorus is generally the limiting nutrient, and the associated risk of algal growth and oxygen depletion. The unit converts nutrient releases and their transport to freshwater into an equivalent mass of phosphorus.
Eutrophication, marine	Fraction of nutrients reaching the marine end compartment (N)	kg N eq	Represents nutrient enrichment of marine waters, where nitrogen is generally the limiting nutrient, and the associated risk of algal blooms and oxygen depletion. The unit expresses releases reaching the sea as an equivalent mass of nitrogen.
Ecotoxicity, freshwater	Comparative Toxic Unit for ecosystems (CTU_e)	CTU_e	Estimates the potentially affected fraction of freshwater species resulting from toxic emissions, accounting for environmental fate and effect. The comparative toxic unit allows different chemicals to be aggregated; it is a modelled potential impact, not a measured pollutant concentration.

Continued on next page

Table 5 – continued from previous page

Impact category	Impact-category indicator	Unit	Interest and interpretation
Land use	Soil quality index	Dimensionless (pt)	Captures soil-quality changes caused by land occupation and transformation. The dimensionless point score aggregates four soil functions: biotic production, erosion resistance, mechanical filtration and groundwater replenishment. It is useful for comparison and aggregation but is not a physical area or a direct biodiversity measure.
Water use	User deprivation potential (deprivation-weighted water consumption)	m^3 water eq of deprived water	Measures the potential deprivation of water for other users and ecosystems, taking regional water scarcity into account. A cubic metre consumed in a water-scarce basin therefore counts more than one consumed in a water-abundant basin; the result is a scarcity-weighted equivalent volume, not the physical withdrawal alone.
Resource use, minerals and metals	Abiotic resource depletion (ADP, ultimate reserves)	kg Sb eq	Indicates the relative depletion of mineral and metal resources from ultimate geological reserves. The unit expresses each resource's scarcity-weighted depletion as the equivalent mass of antimony (Sb), enabling unlike resources to be compared; it does not represent the mass extracted.
Resource use, fossils	Abiotic resource depletion, fossil fuels (ADP-fossil)	MJ	Quantifies depletion of fossil energy carriers (and, in the EF flow list, uranium) according to their energy content. Megajoules provide a common energy basis for comparing different non-renewable fuels; they are not the energy consumed by the assessed product.

Results — embodied impacts of the online system

Table 6: Online data-centre server fleet in 2025

Class	Nodes	Job slots	HS23 [kHS23]	HS23 per slot	Modelled units
A — EB / HLT1 (Gigabyte, GPU)	195	12,800	166.2	12.98	200 nodes
B — HLT2 (Supermicro, Xeon 8592+)	203	51,968	814.0	15.66	200 nodes
C — HLT2 generic (mixed E5 v4 + AMD)	4,901	243,168	2,654.2	10.92	1,225 4-node chassis
Total	5,299	307,936	3,634.4	11.80	1,625 items / 5,300 nodes

Table 7: Annual embodied climate-change impacts of the online system by component

Component	Family	Database	Units	Lifetime [yr]	Per unit [kg CO_2eq]	Annual [t CO_2eq/yr]	Share [%]
Servers class C (HLT2, generic)	Servers	ResilioDB	1,225	10.0	1,993	244.20	41.3
HLT1 storage	Storage	ResilioDB	3,200	10.0	430	137.69	23.3
Servers class B (HLT2)	Servers	ResilioDB	200	10.0	2,267	45.35	7.7
Servers class A (EB/HLT1)	Servers	ResilioDB	200	10.0	2,235	44.70	7.6
HLT2 storage	Storage	ResilioDB	720	10.0	480	34.57	5.9
Containers	Infrastructure	openLCA	7	50.0	188,514	26.39	4.5
Calibration storage	Storage	ResilioDB	600	10.0	331	19.84	3.4
Air handling units	Infrastructure	openLCA	24	20.0	9,471	11.36	1.9
Routers / switches	Network	ResilioDB	318	10.0	354	11.25	1.9
InfiniBand controllers	Network	ResilioDB	600	10.0	65	3.87	0.7
Service storage (NVMe)	Storage	ResilioDB	48	10.0	728	3.50	0.6
Foundations (concrete)	Infrastructure	openLCA	1	50.0	95,360	1.91	0.3
Inverters	Infrastructure	openLCA	144	20.0	210	1.51	0.3
Water pumps	Infrastructure	openLCA	24	20.0	1,236	1.48	0.3
Transformers	Infrastructure	openLCA	2	50.0	30,541	1.22	0.2
Air filters	Infrastructure	openLCA	192	0.5	2	0.87	0.1
Racks	Infrastructure	openLCA	132	50.0	326	0.86	0.1
Reverse osmosis pumps	Infrastructure	openLCA	4	20.0	1,236	0.25	0.0
Cables — Signal	Infrastructure	openLCA	1	50.0	643	0.01	0.0
Cables — ECS	Infrastructure	openLCA	1	50.0	540	0.01	0.0
Cables — TFC	Infrastructure	openLCA	1	50.0	272	0.01	0.0
Total			—	—	—	590.86	100.0

Table 8: Annual embodied impacts of the online system by equipment family — EF 3.1

Impact category	Unit	Total / year	Servers	Storage	Network	Infrastructure
Acidification	mol H+-Eq	3,238	48.7	40.2	2.9	8.2
Climate change	t CO_2eq	591	56.6	33.1	2.6	7.8
Ecotoxicity, freshwater	CTUe	8.48×10^{6}	21.6	66.8	2.5	9.1
Energy resources	MJ	7.80×10^{6}	54.6	36.6	2.5	6.3
Eutrophication, freshwater	kg P-Eq	231	70.1	15.1	2.1	12.7
Eutrophication, marine	kg N-Eq	631	49.6	38.5	2.5	9.4
Eutrophication, terrestrial	mol N-Eq	5,740	51.8	35.5	2.5	10.2
Human toxicity, cancer	CTUh	2.73×10^{-4}	30.4	47.9	2.5	19.3
Human toxicity, non-cancer	CTUh	9.91×10^{-3}	36.3	48.2	2.6	12.9
Ionising radiation	kBq U235-Eq	1.48×10^{6}	19.5	78.4	1.9	0.2
Land use	dimensionless	1.06×10^{6}	55.8	25.6	1.9	16.7
Material resources	kg Sb-Eq	32.5	20.5	68.8	3.6	7.0
Ozone depletion	kg CFC-11-Eq	0.029	26.8	65.7	3.4	4.1
Particulate matter	disease incidence	0.026	43.5	39.2	2.2	15.1
Photochemical ozone	kg NMVOC-Eq	1,791	50.3	37.7	2.5	9.5
Water use	m3 world eq. deprived	200,702	60.8	25.4	2.1	11.7

Table 9: Embodied climate-change impacts of online server processors — EF 3.1

Processor	Class	Node	Launch	Chips per unit	One chip [kg CO_2eq]	BLD [%]	EOL [%]	Logic wafer [%]	Share of the server [%]	Fleet, total [t CO_2eq]
AMD EPYC 7502	A	N7	2019	2	11.42	99.97	0.026	97.9	1.0	4.6
Intel Xeon Platinum 8592+	B	N10	2023	2	139.95	100.00	0.002	99.8	12.3	56.0
Intel Xeon E5-2680 v4	C	N14	2016	8	26.74	99.99	0.011	99.1	10.7	262.1

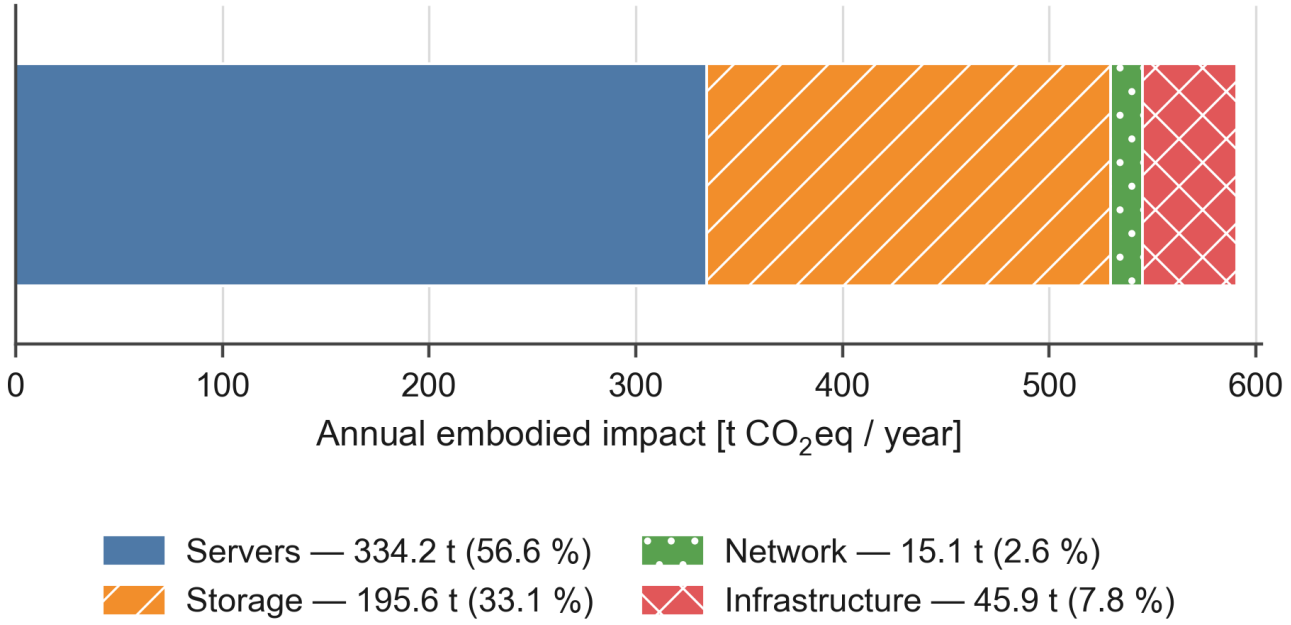


Figure 3: Annual embodied climate-change impact of the online system by equipment family

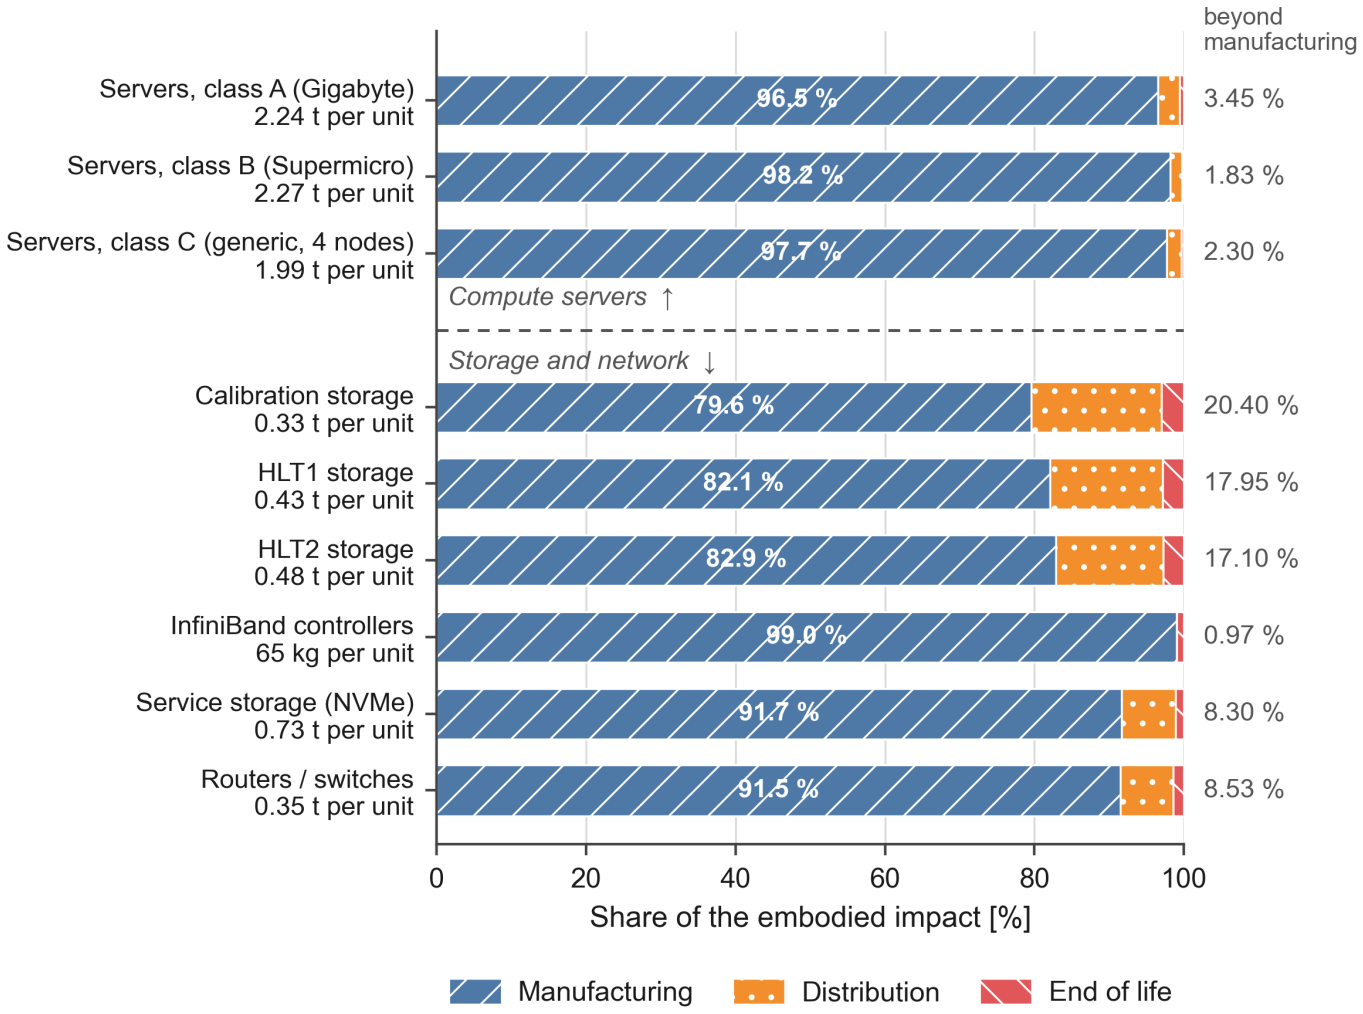


Figure 4: Online IT climate-change impact by life-cycle stage

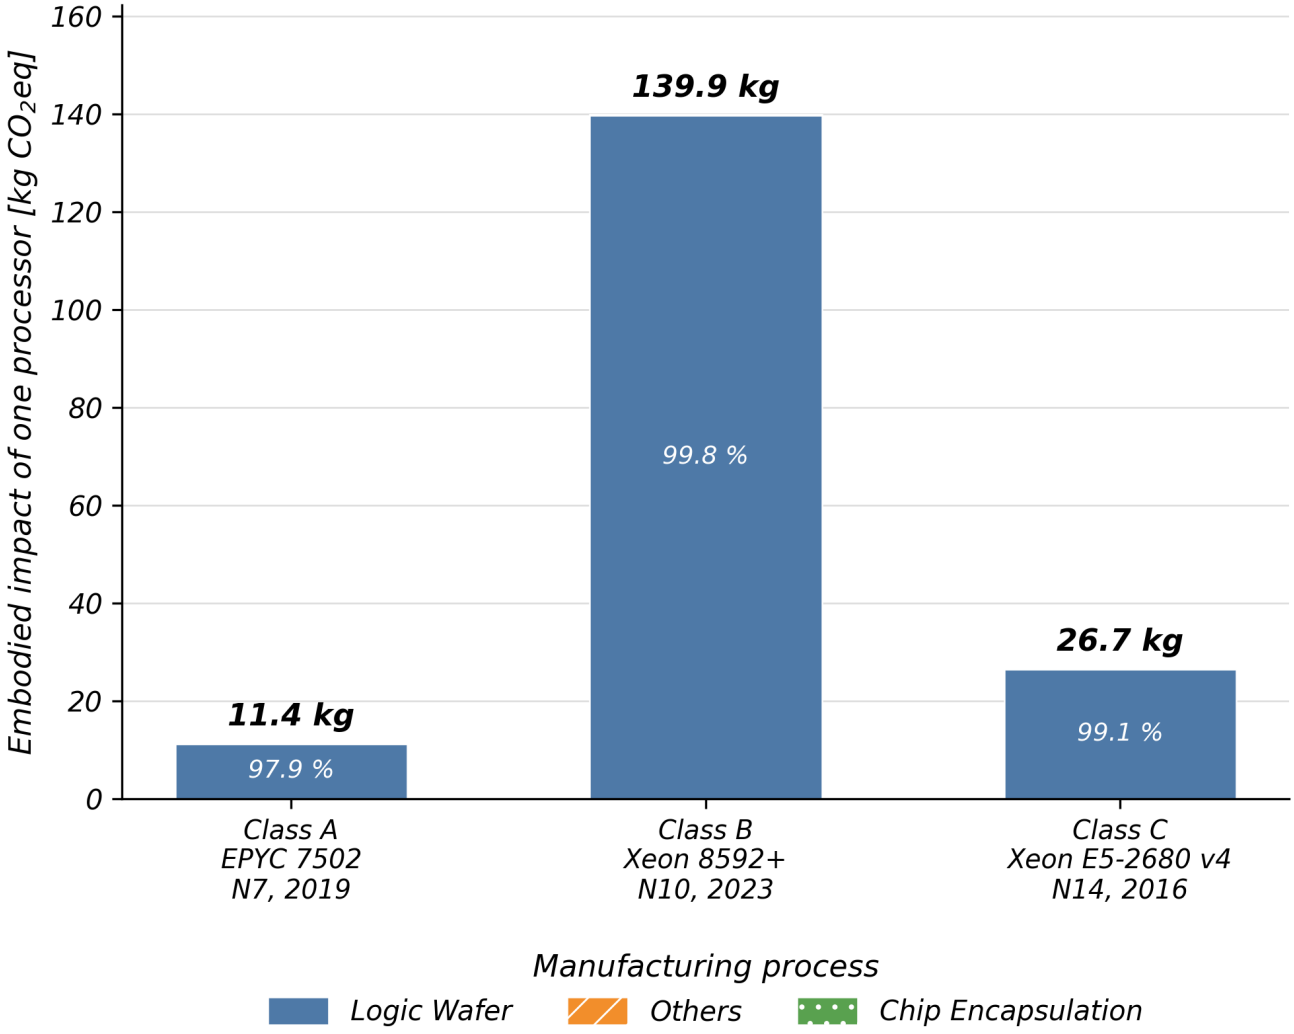


Figure 5: Online server processors by manufacturing process

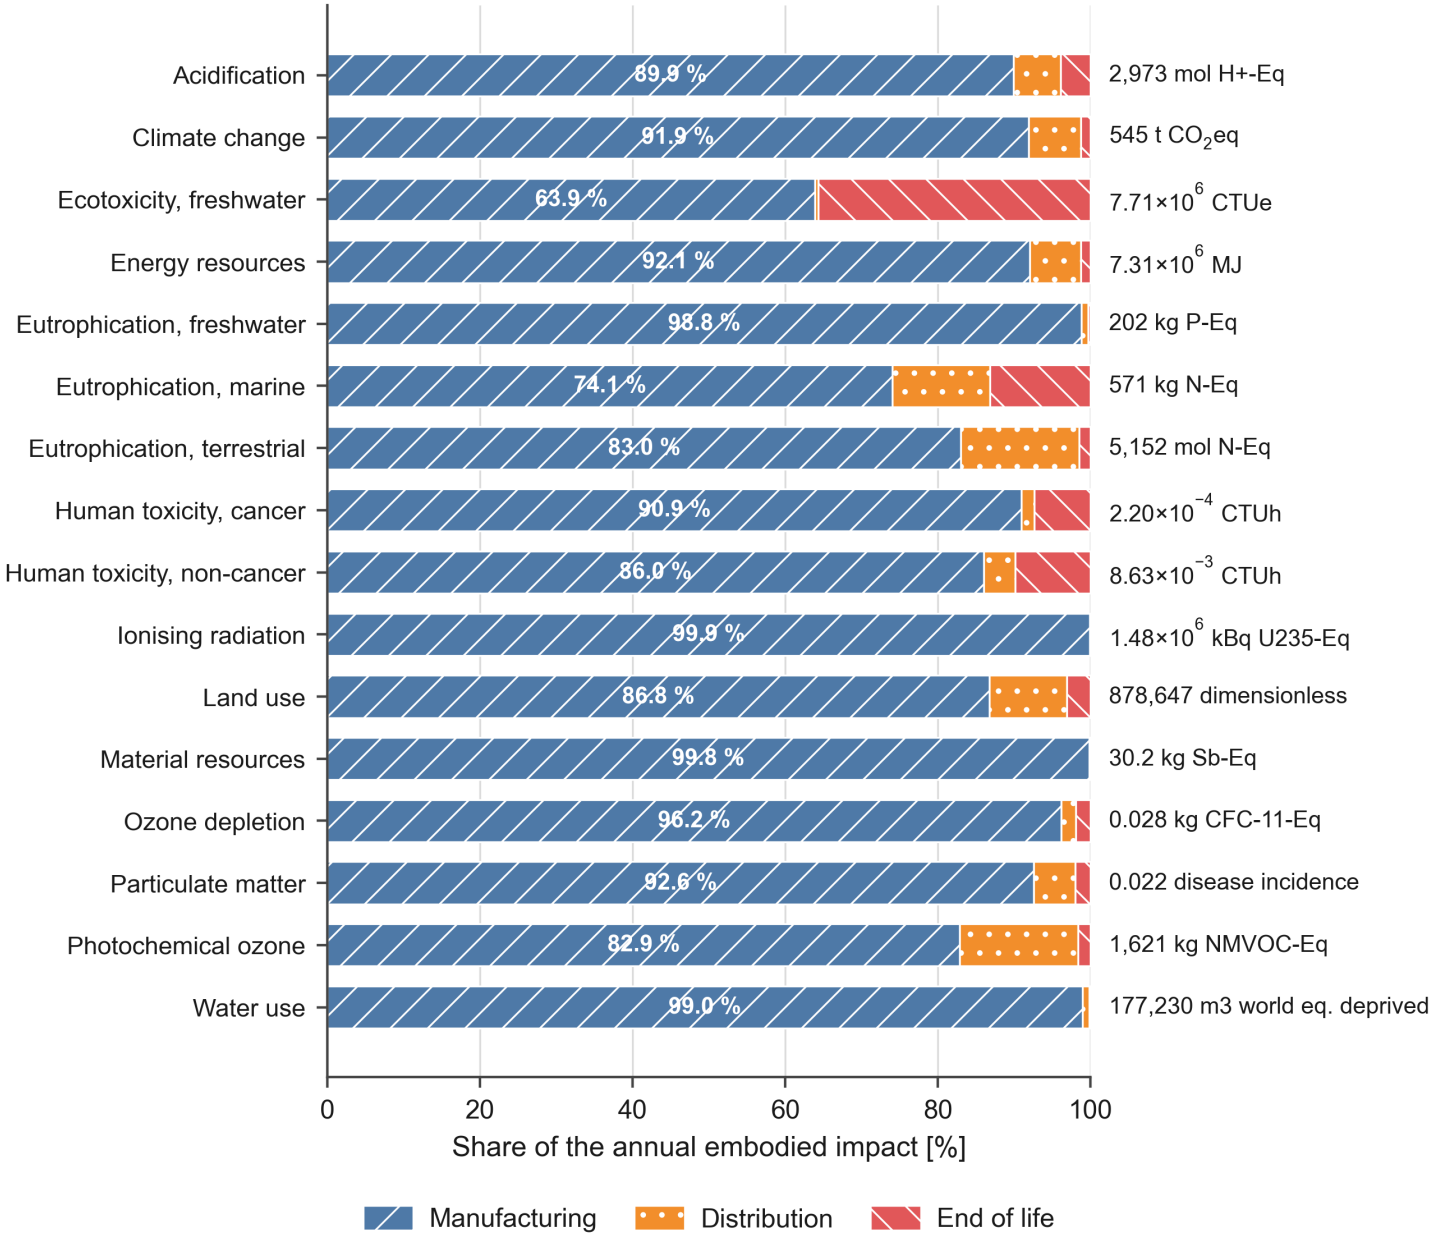


Figure 6: Online IT impacts by life-cycle stage — EF 3.1

Results — operational impacts of the online system

Table 10: Online operational results for 2025

Quantity	Value	Unit
Total metered electricity	11,996	MWh
IT equipment	11,439	MWh
Air handling units	556	MWh
Annual PUE	1.053	–
IT energy E_{IT}	11,388	MWh
Water consumed (adiabatic)	2,516	m^3
Days with water drawn	212	days
Days above 10 m^3	107	days
WUE	0.22	L/kWh
Climate impact — ecoinvent basis	1,453	t CO_2eq
Climate impact — Electricity Maps basis	372	t CO_2eq

Results — offline system

Table 11: Offline activity data and annual embodied impacts in 2025

Quantity	Value	Unit
Wall-clock time consumed by LHCb	1.10	10^9 logical core·h
Server-equivalents occupied year-round	490	servers
Fleet generational position	14.8	HS23 / job slot
Disk capacity provisioned (annual mean)	146.3	PB
Hard drives provisioned (16 TB reference)	9,144	drives
Embodied impact — compute	111.0	t CO_2eq / yr
Embodied impact — disk storage	416.2	t CO_2eq / yr
Embodied impact — total	527.2	t CO_2eq / yr

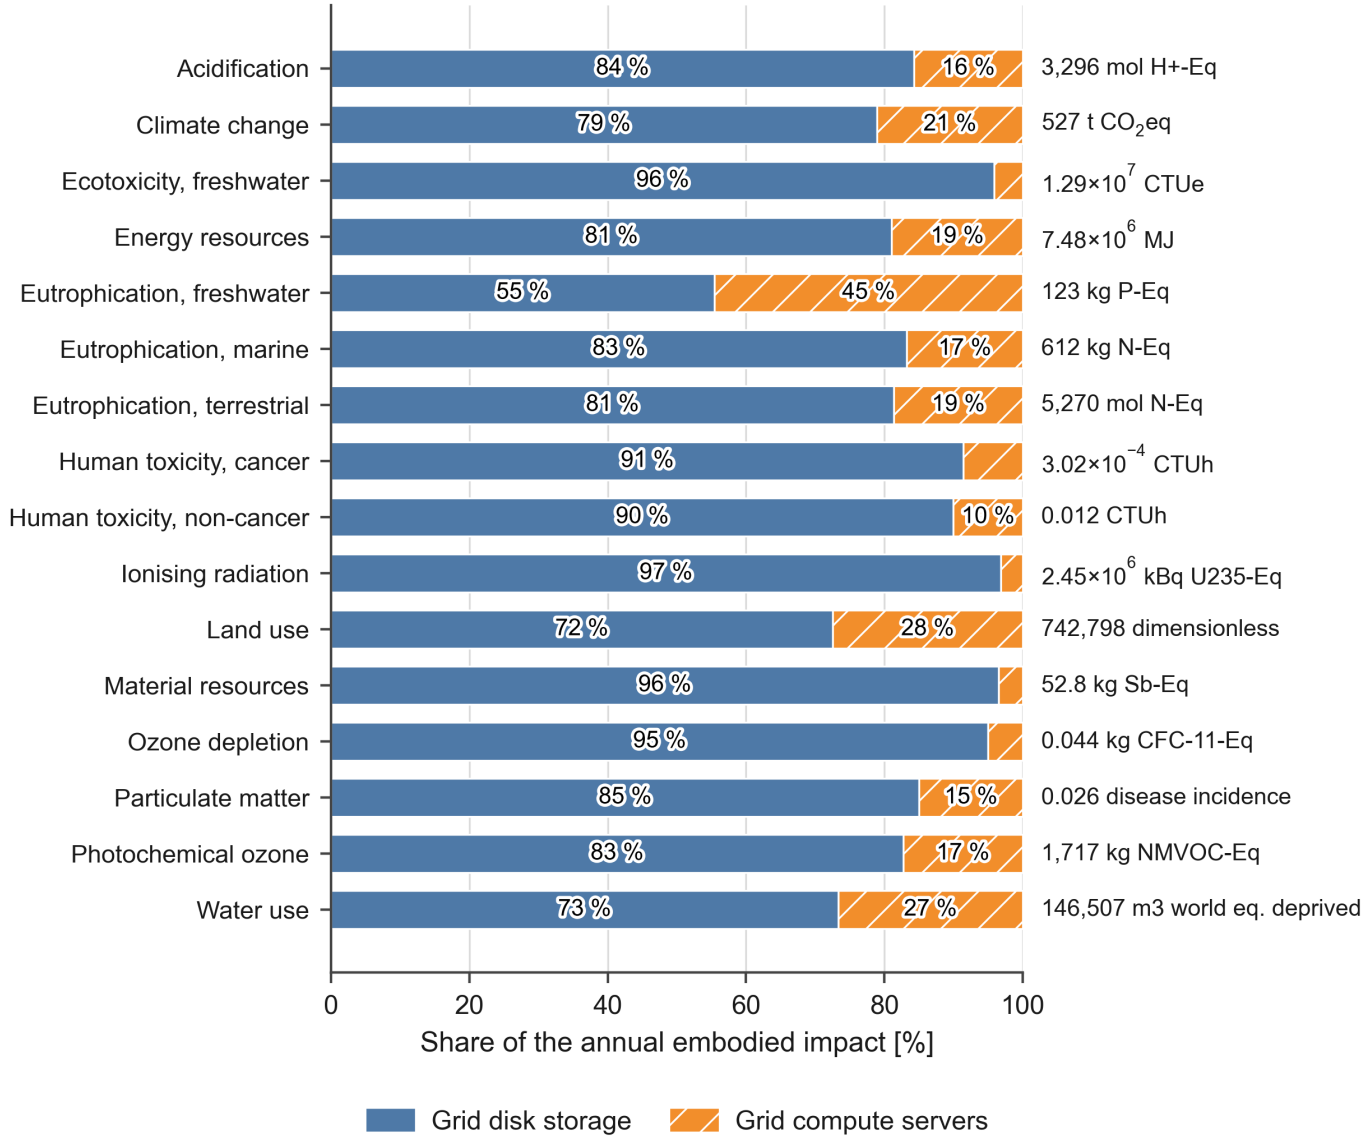


Figure 7: Annual embodied impacts of the offline system — EF 3.1

Table 12: Offline embodied climate-change impacts by scenario

Scenario	Reference drive [TB]	Storage intensity [kg CO_2eq/TB/yr]	Embodied [t CO_2eq/yr]	Deviation [%]
Central (class B, 16 TB, provisioned)	16	2.845	527.2	+0.0
Reference class C (bracket, ~2016 hardware)	16	2.845	542.4	+2.9
Storage allocated on used capacity	16	2.845	387.3	−26.5
HDD reference 12 TB	12	3.378	605.3	+14.8
HDD reference 14 TB	14	3.074	560.7	+6.3
HDD reference 18 TB	18	2.667	501.2	−4.9
HDD reference 20 TB	20	2.525	480.4	−8.9
HDD reference 22 TB	22	2.409	463.4	−12.1
Envelope — lower bound	—	—	387.3	−26.5
Envelope — upper bound	—	—	605.3	+14.8

Table 13: Median HEPScore23 per job slot by processor launch year

Launch year	CPU models [n]	Median HS23/job slot
2017	4	10.6
2019	21	14.8
2021	27	16.7
2022	23	19.7
2023	17	14.9
2024	24	21.8

Table 14: Offline operational results by country in 2025 — Electricity Maps

Country	Sites	Wall-clock [10^6 core·h]	Electricity [MWh]	CI [g CO_2eq/kWh]	Compute [t CO_2eq]	Disk [t CO_2eq]	Total [t CO_2eq]
GB	15	337.4	2,691	176	473.5	21.3	494.8
IT	4	114.1	910	286	260.2	21.2	281.4
PL	3	49.9	398	565	225.0	8.6	233.6
DE	2	69.5	554	342	189.5	25.6	215.1
CN	1	22.3	178	485	86.1	34.2	120.3
FR	6	378.1	3,016	31	93.5	9.8	103.3
NL	2	26.2	209	253	52.8	7.3	60.0
ES	1	25.1	200	136	27.2	3.3	30.5
BR	1	32.1	256	98	25.1	0.0	25.1
CH	1	33.5	268	50	13.4	1.4	14.8
US	1	4.0	32	414	13.1	0.0	13.1
CZ	1	2.8	22	439	9.7	0.0	9.7
RO	1	3.1	25	300	7.5	0.6	8.0
Total	39	1,098.1	8,757	—	1,476.6	133.2	1,609.8

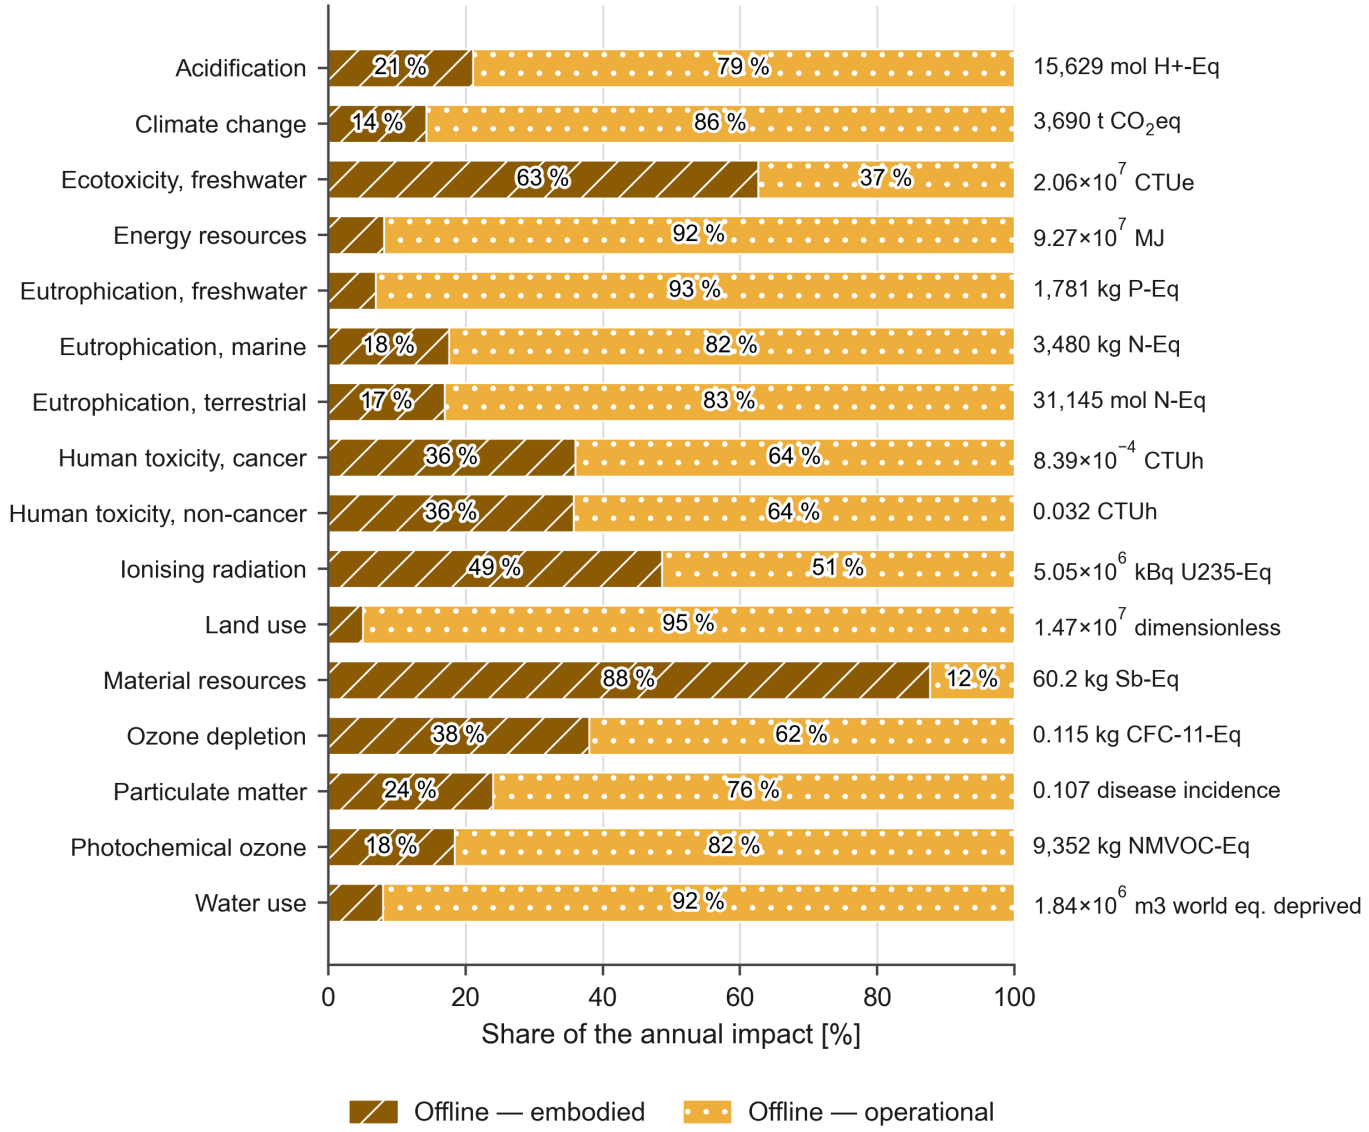


Figure 8: Annual impacts of the offline system — EF 3.1

Table 15: Offline operational climate impact by carbon-intensity source

Source	System boundary	FR [g CO_2eq/kWh]	Climate [t CO_2eq/yr]	vs Electricity Maps [%]
ElectricityMaps	life cycle	31.0	1,610	+0.0
Ember	power sector emissions	41.4	1,773	+10.1
CaDI	life cycle + T&D + well-to-tank	63.0	2,215	+37.6

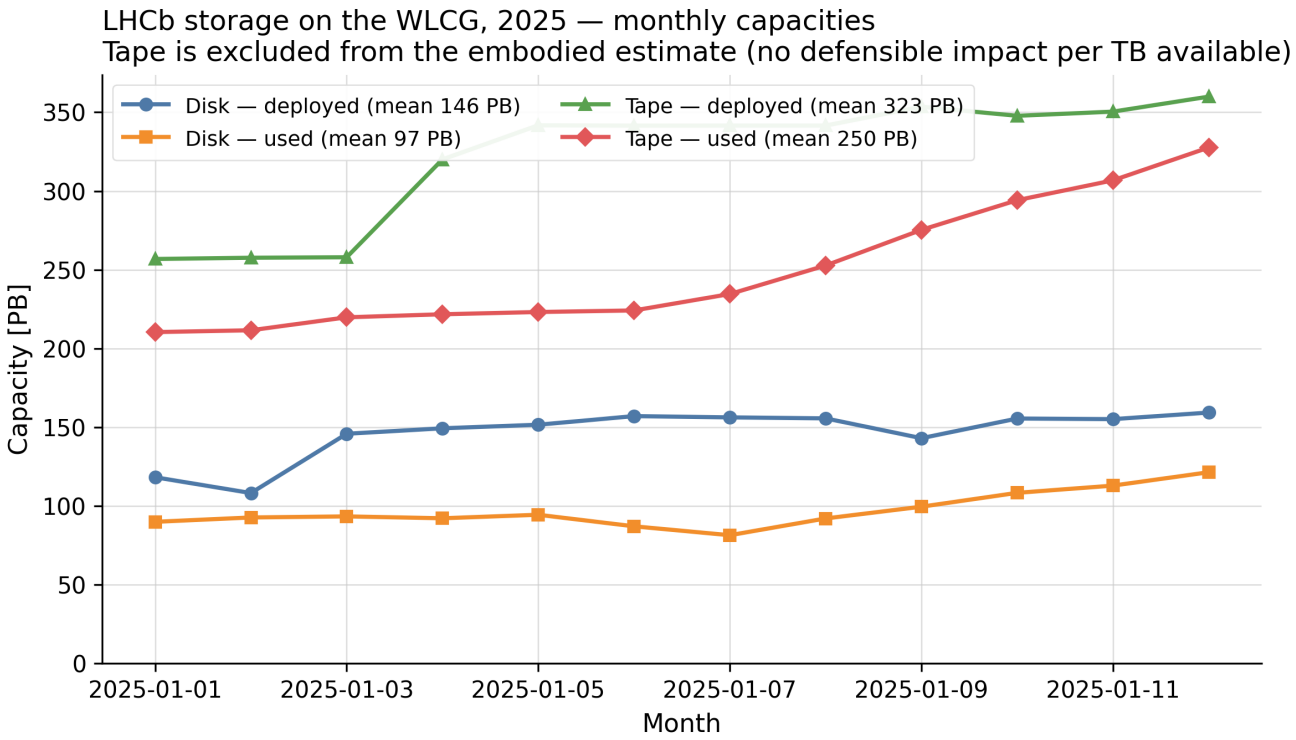


Figure 9: Monthly LHCb grid storage in 2025

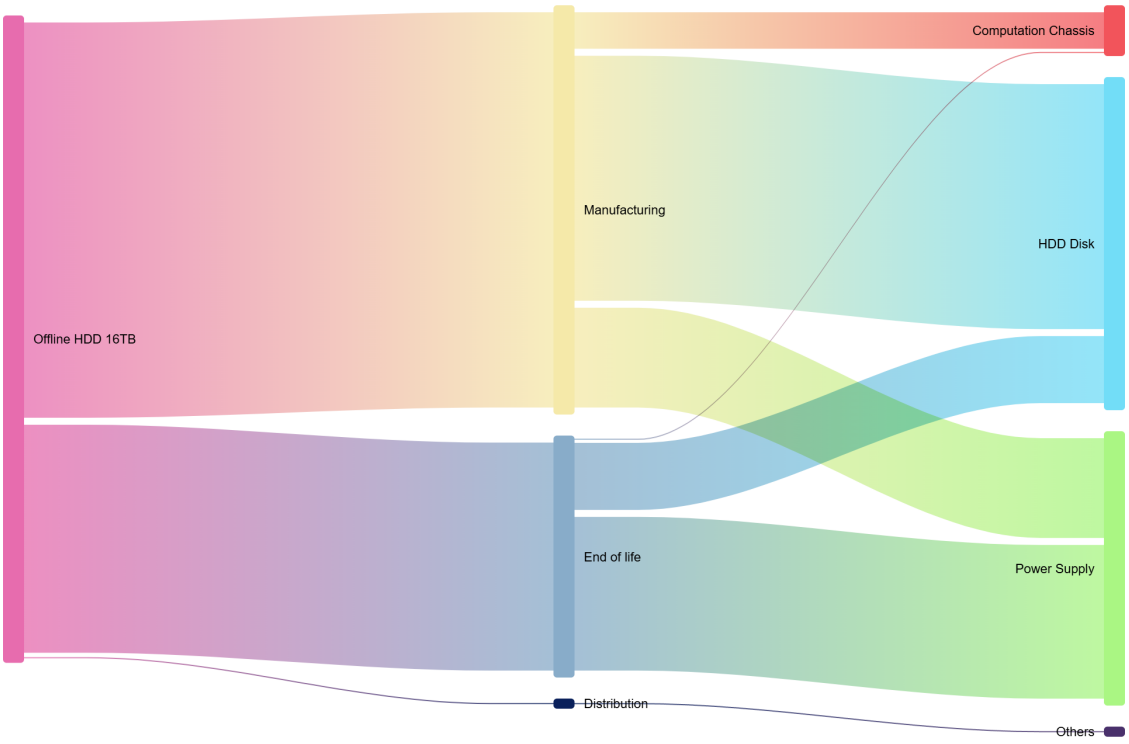


Figure 10: Freshwater-ecotoxicity impact of a 16 TB grid HDD — ResilioDB

Results — multi-impact profile of the complete system

Table 16: Annual impacts by system and nature — EF 3.1, ecoinvent 3.12 APOS

System	Impact category	Unit	Embodied	Operational	Total
Online	Acidification	mol H+-Eq	3,238	5,539	8,777
Online	Climate change	t CO_2eq	591	1,454	2,044
Online	Ecotoxicity, freshwater	CTUe	8.48×10^{6}	3.28×10^{6}	1.18×10^{7}
Online	Energy resources	MJ	7.80×10^{6}	1.28×10^{8}	1.36×10^{8}
Online	Eutrophication, freshwater	kg P-Eq	231	462	693
Online	Eutrophication, marine	kg N-Eq	631	1,587	2,218
Online	Eutrophication, terrestrial	mol N-Eq	5,740	13,085	18,825
Online	Human toxicity, cancer	CTUh	2.73×10^{-4}	3.62×10^{-4}	6.35×10^{-4}
Online	Human toxicity, non-cancer	CTUh	9.91×10^{-3}	0.012	0.022
Online	Ionising radiation	kBq U235-Eq	1.48×10^{6}	5.46×10^{6}	6.94×10^{6}
Online	Land use	dimensionless	1.06×10^{6}	6.71×10^{6}	7.77×10^{6}
Online	Material resources	kg Sb-Eq	32.5	7.45	40.0
Online	Ozone depletion	kg CFC-11-Eq	0.029	0.038	0.066
Online	Particulate matter	disease incidence	0.026	0.059	0.085
Online	Photochemical ozone	kg NMVOC-Eq	1,791	4,273	6,065
Online	Water use	m3 world eq. deprived	200,702	1.55×10^{6}	1.75×10^{6}
Offline	Acidification	mol H+-Eq	3,296	12,333	15,629
Offline	Climate change	t CO_2eq	527	3,162	3,690
Offline	Ecotoxicity, freshwater	CTUe	1.29×10^{7}	7.69×10^{6}	2.06×10^{7}
Offline	Energy resources	MJ	7.48×10^{6}	8.52×10^{7}	9.27×10^{7}
Offline	Eutrophication, freshwater	kg P-Eq	123	1,658	1,781
Offline	Eutrophication, marine	kg N-Eq	612	2,868	3,480
Offline	Eutrophication, terrestrial	mol N-Eq	5,270	25,875	31,145
Offline	Human toxicity, cancer	CTUh	3.02×10^{-4}	5.37×10^{-4}	8.39×10^{-4}
Offline	Human toxicity, non-cancer	CTUh	0.012	0.021	0.032
Offline	Ionising radiation	kBq U235-Eq	2.45×10^{6}	2.59×10^{6}	5.05×10^{6}
Offline	Land use	dimensionless	742,798	1.39×10^{7}	1.47×10^{7}
Offline	Material resources	kg Sb-Eq	52.8	7.44	60.2
Offline	Ozone depletion	kg CFC-11-Eq	0.044	0.071	0.115
Offline	Particulate matter	disease incidence	0.026	0.081	0.107
Offline	Photochemical ozone	kg NMVOC-Eq	1,717	7,635	9,352
Offline	Water use	m3 world eq. deprived	146,507	1.69×10^{6}	1.84×10^{6}

Table 17: Monte Carlo mean annual impacts of the complete system — EF 3.1

Impact category	Unit	Online — embodied	Online — operational	Offline — embodied	Offline — operational	Total	Total P5–P95
Acidification	mol H+-Eq	3,502	5,446	3,617	12,259	24,825	20,816–28,881
Climate change	t CO_2eq	637	1,425	578	3,143	5,784	4,810–6,750
Ecotoxicity, freshwater	CTUe	9.10×10^{6}	3.23×10^{6}	1.41×10^{7}	7.60×10^{6}	3.41×10^{7}	2.84×10^{7}–4.15×10^{7}
Energy resources	MJ	8.43×10^{6}	1.26×10^{8}	8.20×10^{6}	8.45×10^{7}	2.27×10^{8}	1.65×10^{8}–3.10×10^{8}
Eutrophication, freshwater	kg P-Eq	249	439	135	1,655	2,478	1,416–4,180
Eutrophication, marine	kg N-Eq	677	1,561	671	2,849	5,758	4,797–6,777
Eutrophication, terrestrial	mol N-Eq	6,171	12,947	5,781	25,692	50,591	42,087–59,716
Human toxicity, cancer	CTUh	2.86×10^{-4}	3.64×10^{-4}	3.31×10^{-4}	5.39×10^{-4}	1.52×10^{-3}	1.27×10^{-3}–1.81×10^{-3}
Human toxicity, non-cancer	CTUh	0.011	0.012	0.013	0.021	0.056	0.045–0.068
Ionising radiation	kBq U235-Eq	1.63×10^{6}	5.01×10^{6}	2.69×10^{6}	2.46×10^{6}	1.18×10^{7}	6.87×10^{6}–2.24×10^{7}
Land use	dimensionless	1.12×10^{6}	6.56×10^{6}	814,705	1.39×10^{7}	2.24×10^{7}	1.75×10^{7}–2.78×10^{7}
Material resources	kg Sb-Eq	35.2	7.25	57.9	7.44	108	88.0–136
Ozone depletion	kg CFC-11-Eq	0.031	0.037	0.048	0.071	0.187	0.156–0.224
Particulate matter	disease incidence	0.028	0.057	0.028	0.079	0.192	0.150–0.251
Photochemical ozone	kg NMVOC-Eq	1,926	4,205	1,884	7,569	15,584	13,028–18,316
Water use	m3 world eq. deprived	215,158	1.53×10^{6}	160,695	1.68×10^{6}	3.60×10^{6}	2.87×10^{6}–4.43×10^{6}

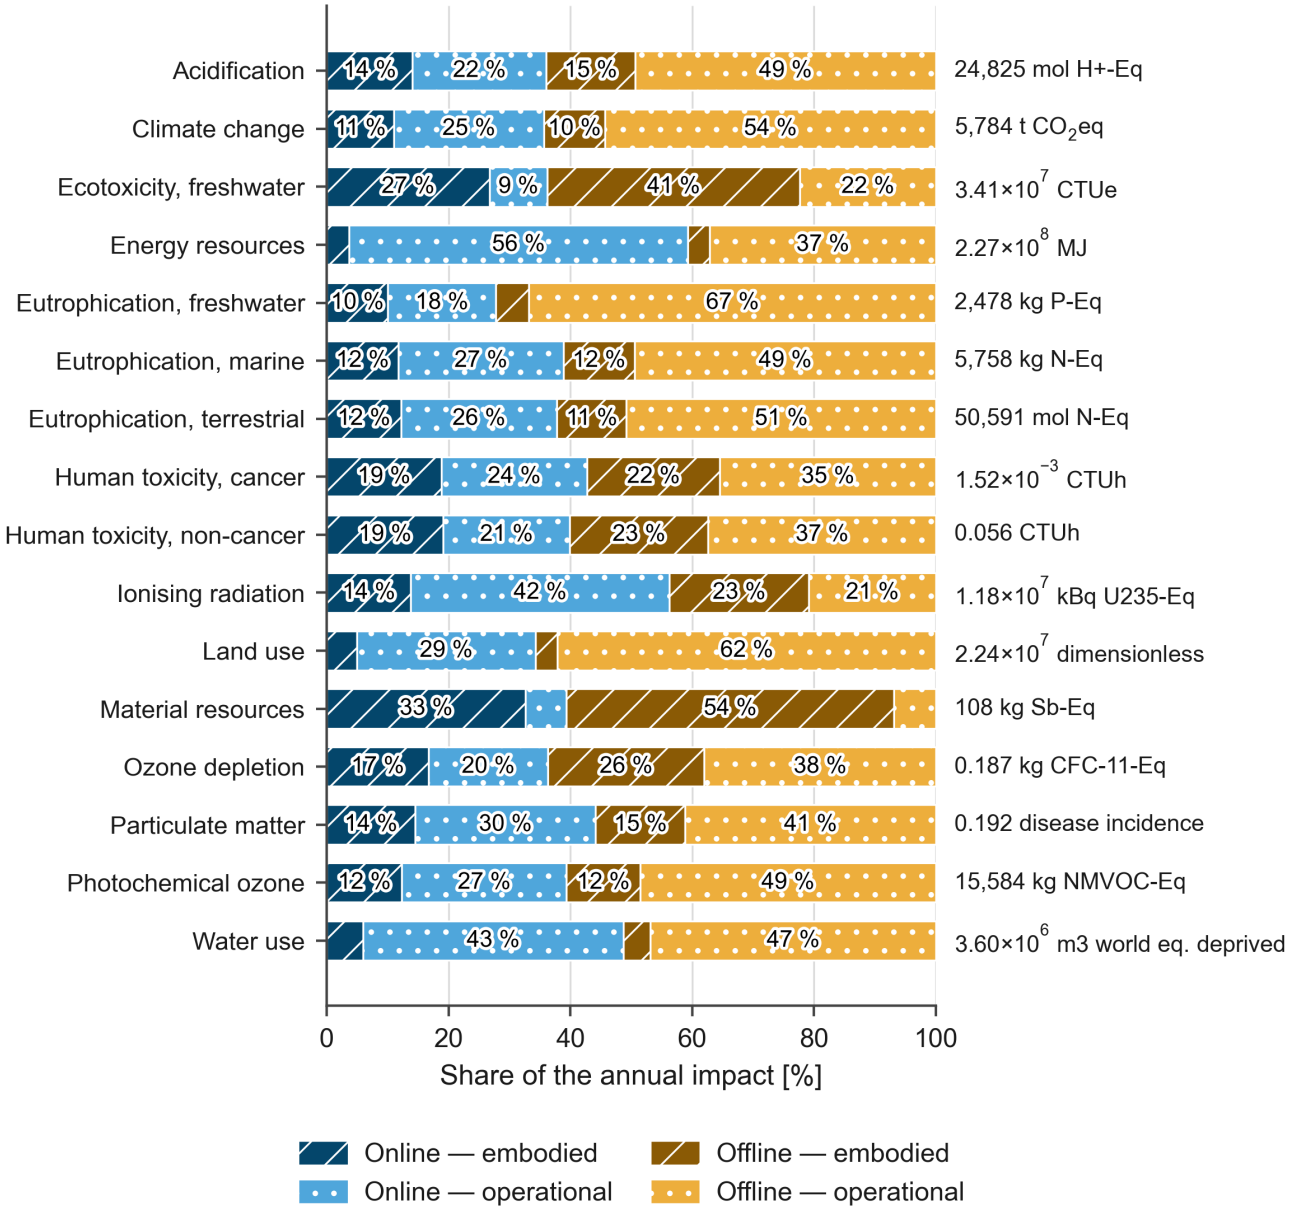


Figure 11: Complete-system impact profile — EF 3.1, ecoinvent 3.12 APOS

Results — climate-change balance and grid geography

Table 18: Climate-change balance in 2025 — Electricity Maps

System	Nature	Post	t CO_2eq / year	Share [%]
Online	Embodied	Servers	334.2	10.8
Online	Embodied	Storage	195.6	6.3
Online	Embodied	Network	15.1	0.5
Online	Embodied	Infrastructure	45.9	1.5
Online	Operational	Electricity	371.9	12.0
Offline	Embodied	Servers	111.0	3.6
Offline	Embodied	Storage	416.2	13.4
Offline	Operational	Electricity (compute)	1,476.6	47.6
Offline	Operational	Electricity (disk)	133.2	4.3
Total			3,099.8	100.0

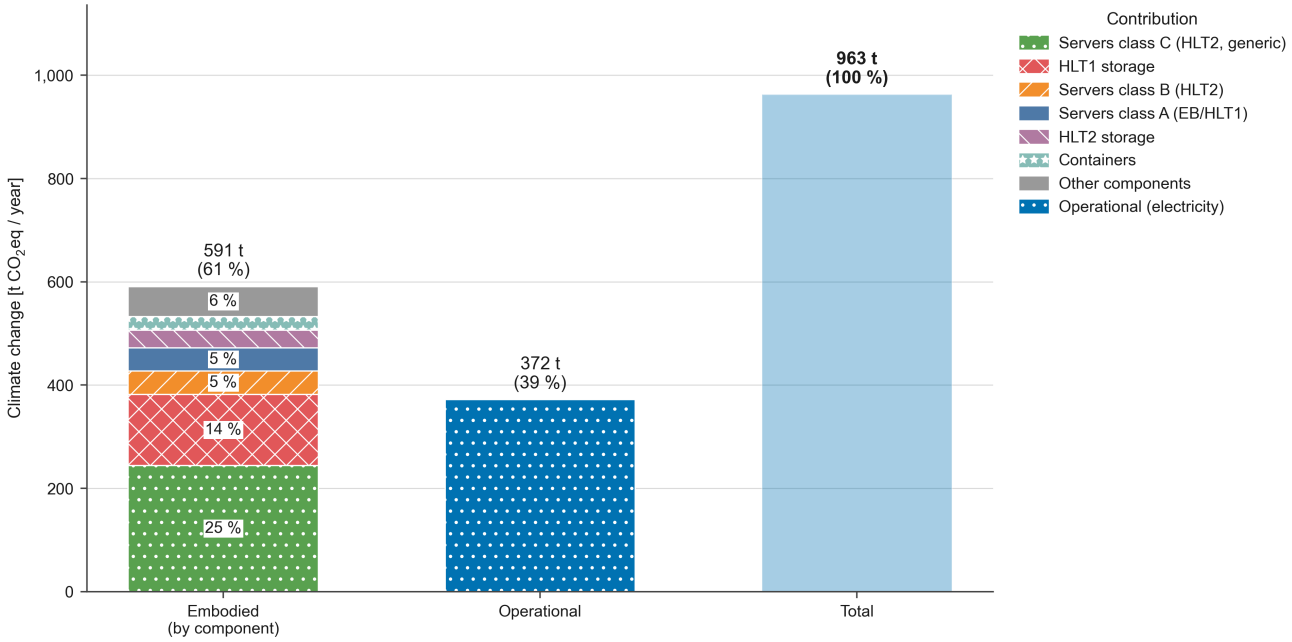


Figure 12: Online climate-change impacts by component — Electricity Maps

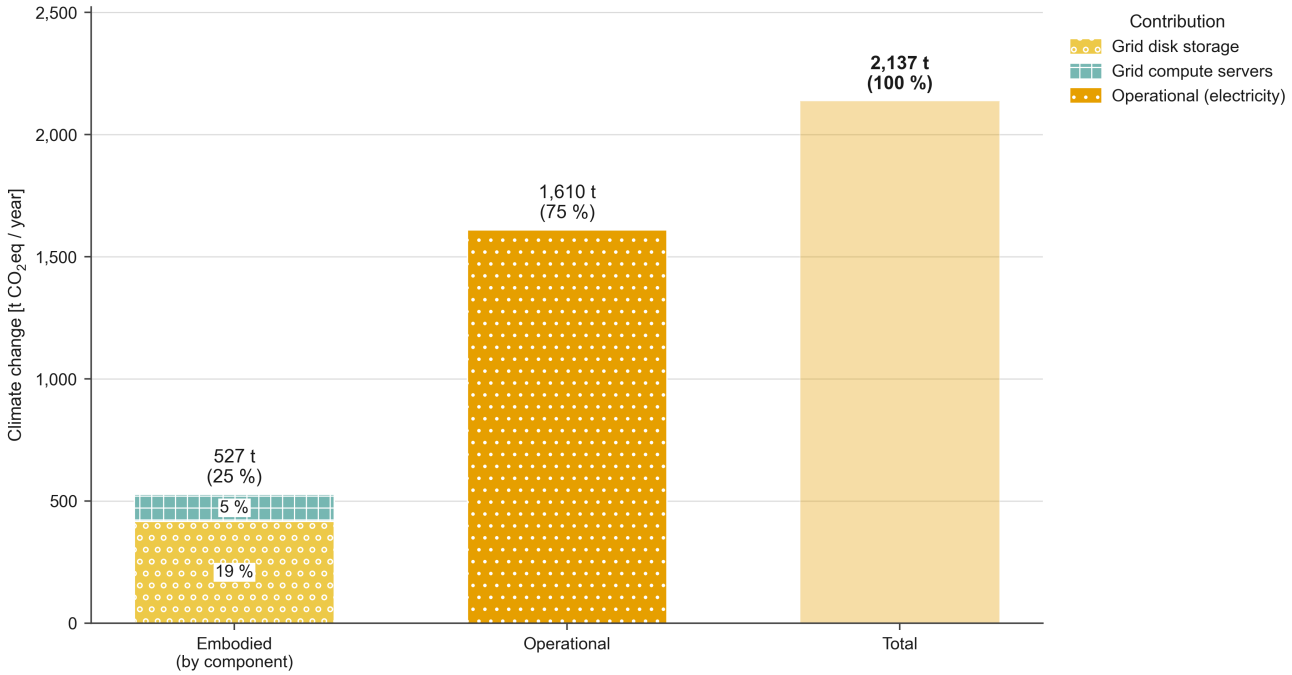


Figure 13: Offline climate-change impacts by component — Electricity Maps

Table 19: Offline computing activity and climate impact by host country — Electricity Maps

Country	Sites	Work [10^9 HS23·h]	(topology) [10^9]	Electricity [MWh]	Mix [g CO_2eq/kWh]	Impact [t CO_2eq/yr]	Share work [%]	Share CO_2 [%]	Intensity [g CO_2eq/kHS23·h]
France	6	5.098	1.438	3,015.5	31	93.5	31.9	6.4	18.3
United Kingdom	15	4.488	4.488	2,690.6	176	473.5	28.1	32.3	105.5
Italy	4	2.456	2.456	909.7	286	260.2	15.4	17.7	105.9
Germany	2	1.076	1.076	554.2	342	189.5	6.7	12.9	176.1
Poland	3	0.708	0.708	398.2	565	225.0	4.4	15.3	317.9
Brazil	1	0.524	0.524	256.2	98	25.1	3.3	1.7	47.9
China	1	0.492	0.492	177.5	485	86.1	3.1	5.9	175.1
Netherlands	2	0.447	0.447	208.6	253	52.8	2.8	3.6	118.1
Spain	1	0.329	0.329	200.3	136	27.2	2.1	1.9	82.8
Switzerland	1	0.311	3.971	267.6	50	13.4	1.9	0.9	43.0
United States	1	0.036	0.036	31.7	414	13.1	0.2	0.9	361.8
Romania	1	0.033	0.033	24.9	300	7.5	0.2	0.5	228.0
Total (mapped sites)	38	15.999	15.999	8,734.9	—	1,466.9	100.0	100.0	91.7

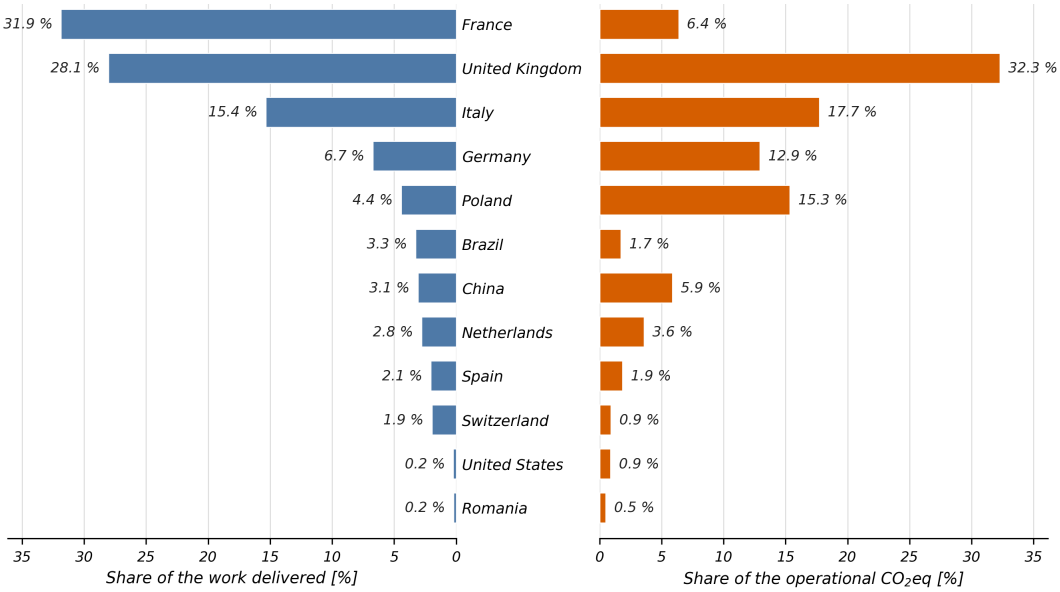


Figure 14: Offline computing and climate-impact shares — Electricity Maps

Table 20: Offline operational climate impact by host country — Electricity Maps

Country	Sites	Compute [MWh]	Disk [MWh]	Mix [g CO_2eq/kWh]	Compute [t CO_2eq]	Disk [t CO_2eq]	Total [t CO_2eq]	Share of total [%]	Storage in country [%]
United Kingdom	15	2,690.5	120.9	176	473.5	21.3	494.8	30.7	4.3
Italy	4	909.7	74.2	286	260.2	21.2	281.4	17.5	7.5
Poland	3	398.2	15.2	565	225.0	8.6	233.6	14.5	3.7
Germany	2	554.2	74.9	342	189.5	25.6	215.1	13.4	11.9
China	1	177.5	70.5	485	86.1	34.2	120.3	7.5	28.4
France	6	3,015.5	316.8	31	93.5	9.8	103.3	6.4	9.5
Netherlands	2	208.6	28.7	253	52.8	7.3	60.0	3.7	12.1
Spain	1	200.3	24.2	136	27.2	3.3	30.5	1.9	10.8
Brazil	1	256.2	0.0	98	25.1	0.0	25.1	1.6	0.0
Switzerland	1	267.6	27.6	50	13.4	1.4	14.8	0.9	9.4
United States	1	31.7	0.0	414	13.1	0.0	13.1	0.8	0.0
Czechia	1	22.1	0.0	439	9.7	0.0	9.7	0.6	0.0
Romania	1	24.9	1.9	300	7.5	0.6	8.0	0.5	7.3
Total	39	8,757.0	754.9	—	1,476.6	133.2	1,609.8	100.0	8.3

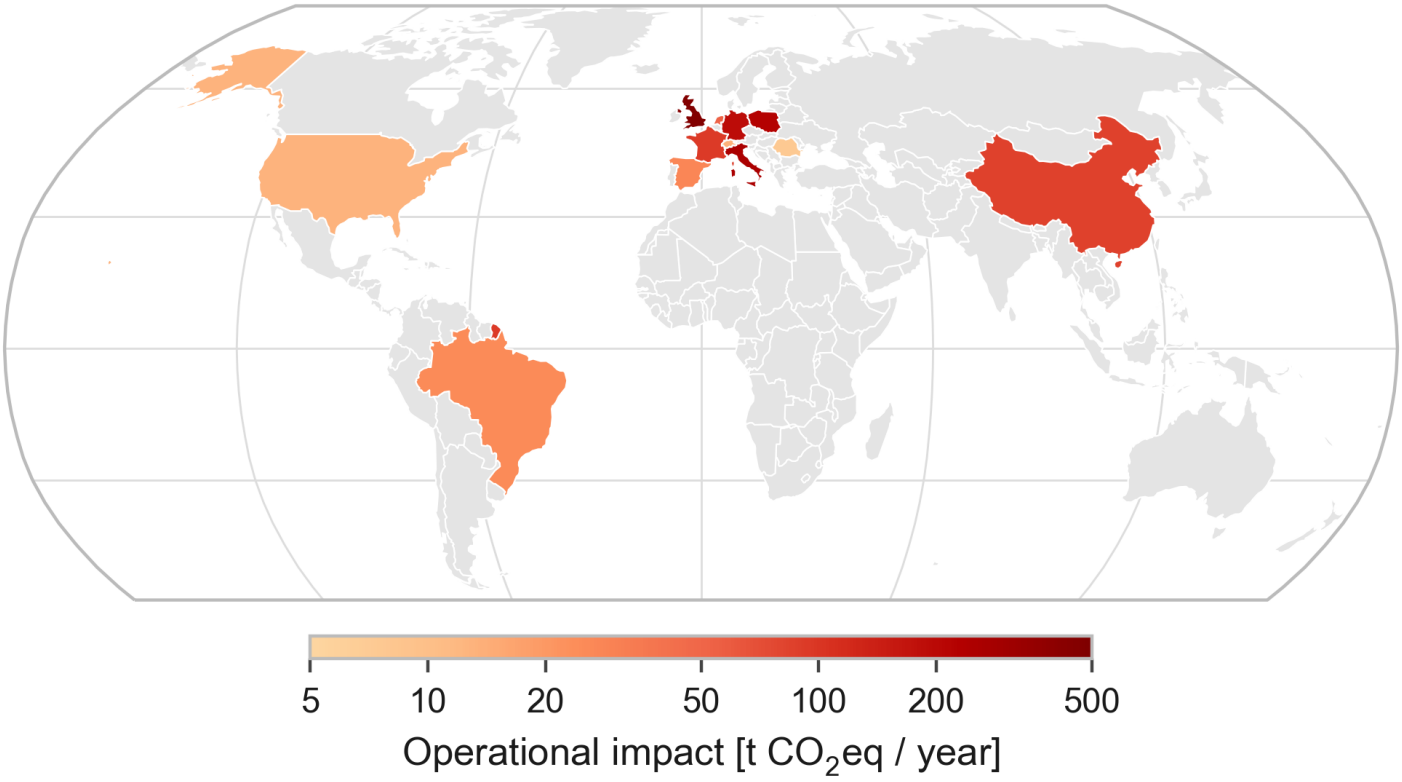


Figure 15: Offline computing climate-change impact by host country — Electricity Maps

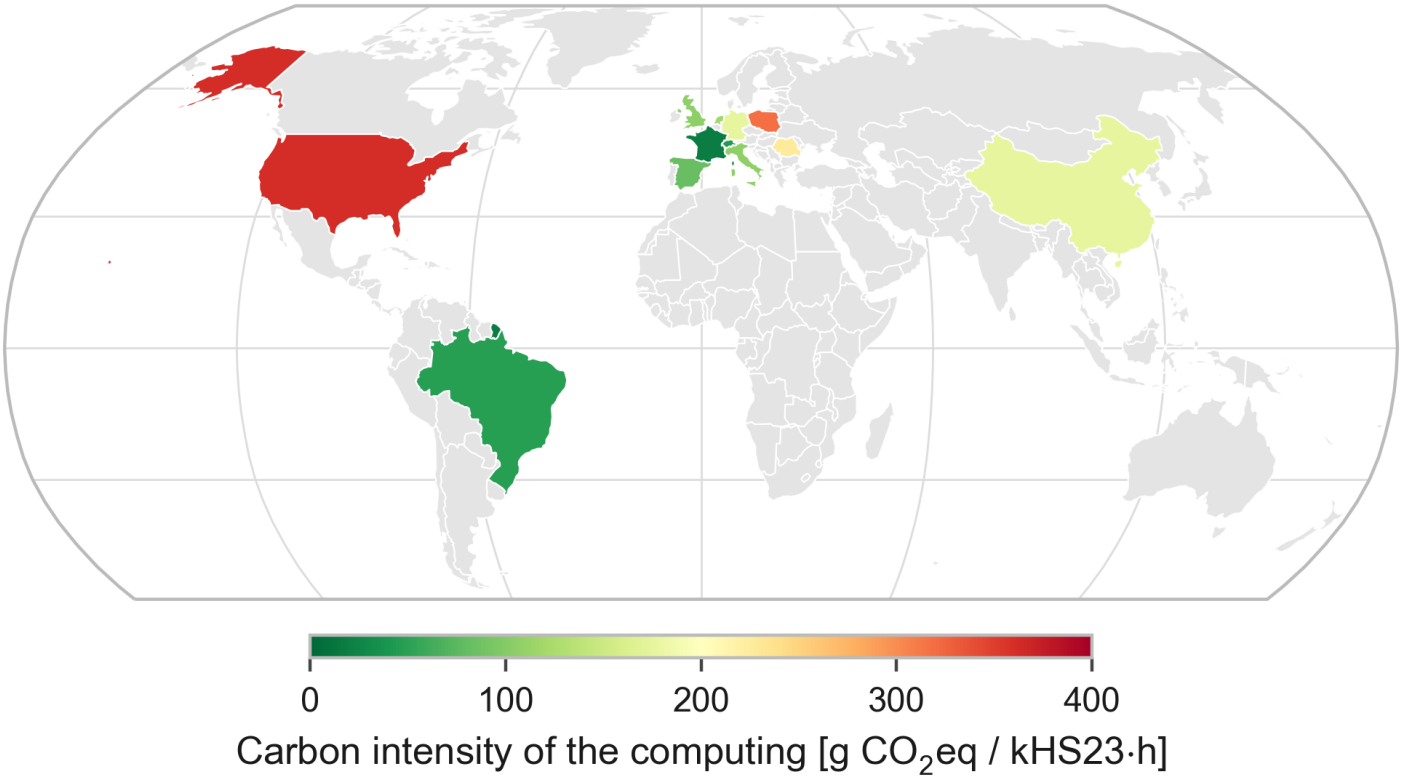


Figure 16: Grid-computing carbon intensity by host country

Results — uncertainty analysis

Table 21: Monte Carlo climate results in 2025 — Electricity Maps

Scope	Nominal [t]	MC mean [t]	MC median [t]	P5–P95 [t]	Mean vs nominal [%]
Full system	3,100	3,182	3,177	2,644–3,729	2.7
Online (Point 8)	963	1,009	991	892–1,188	4.8
Offline (WLCG)	2,137	2,173	2,174	1,685–2,654	1.7

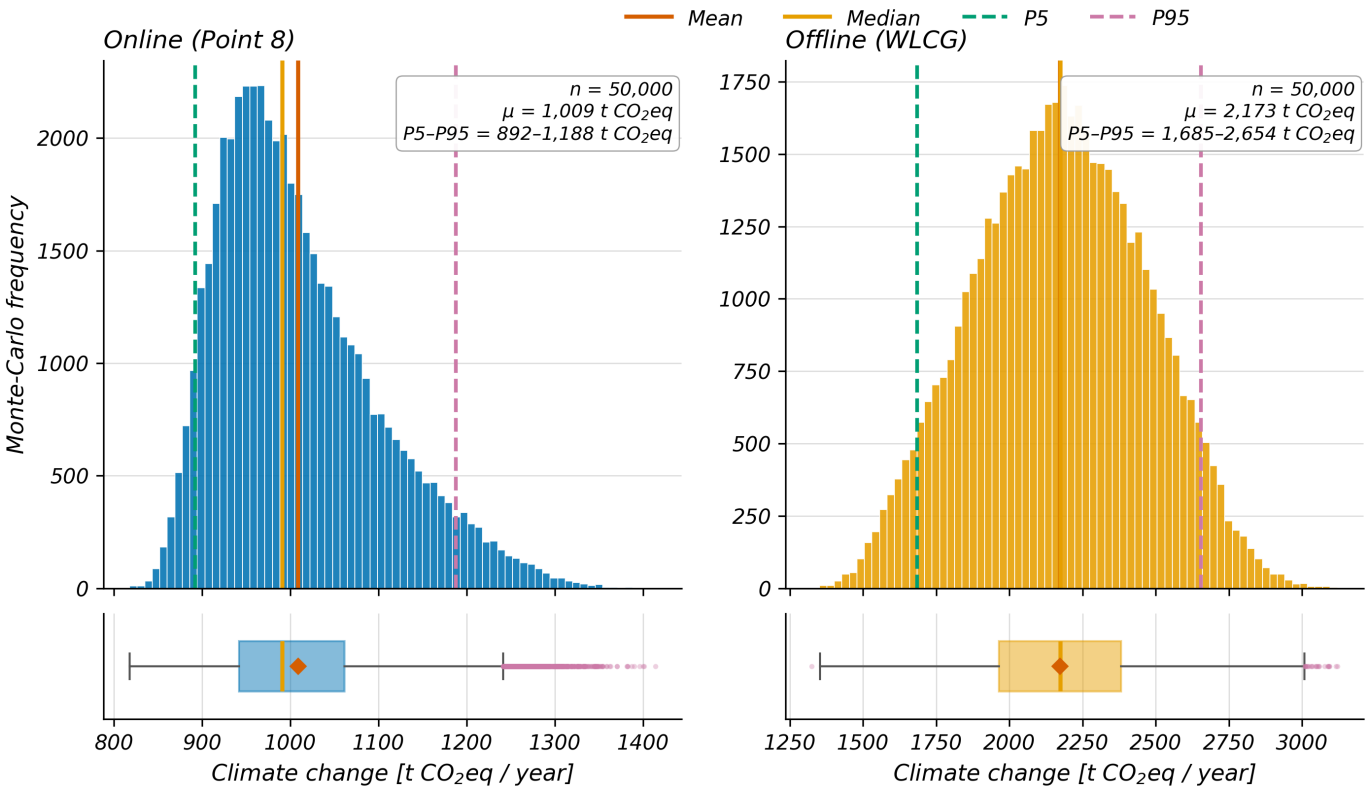


Figure 17: Online and offline Monte Carlo distributions — Electricity Maps

Table 22: First-order variance shares and confidence intervals

Result	Group	Share [%]	IC90 [%]
Headline (Electricity Maps)	G1 — dataset quality	1.53	1.53–1.54
Headline (Electricity Maps)	G2 — representativeness	70.47	70.46–70.49
Headline (Electricity Maps)	G3 — input data	0.59	0.59–0.59
Headline (Electricity Maps)	G4 — lifespan	26.57	26.56–26.57
Headline (Electricity Maps)	Interaction residual	0.84	0.81–0.85
Multi-impact climate row (ecoinvent)	G1 — dataset quality	5.31	4.95–5.80
Multi-impact climate row (ecoinvent)	G2 — representativeness	85.69	85.19–86.66
Multi-impact climate row (ecoinvent)	G3 — input data	0.61	0.61–0.62
Multi-impact climate row (ecoinvent)	G4 — lifespan	8.37	8.31–8.47
Multi-impact climate row (ecoinvent)	Interaction residual	0.03	−0.97–0.37

Table 23: Monte Carlo sensitivity scenarios

Axis	Variant	Mean [t]	Median [t]	P5 [t]	P95 [t]	Delta mean [%]
hdd_capacity_tb	16 TB	3,182.0	3,177.1	2,644.0	3,729.2	0.0
hdd_capacity_tb	12 TB	3,267.6	3,261.6	2,719.2	3,831.7	2.7
hdd_capacity_tb	14 TB	3,218.7	3,213.3	2,676.2	3,773.3	1.2
hdd_capacity_tb	18 TB	3,153.4	3,148.5	2,617.9	3,695.8	−0.9
hdd_capacity_tb	20 TB	3,130.6	3,125.9	2,597.5	3,668.2	−1.6
hdd_capacity_tb	22 TB	3,111.9	3,107.4	2,580.5	3,647.0	−2.2
offline_reference_class	servers_B	3,182.0	3,177.1	2,644.0	3,729.2	0.0
offline_reference_class	servers_C	3,198.7	3,192.7	2,656.8	3,750.2	0.5
storage_allocation	provisioned_pb	3,182.0	3,177.1	2,644.0	3,729.2	0.0
storage_allocation	used_pb	3,136.1	3,131.1	2,597.9	3,683.6	−1.4
ci_source	ElectricityMaps	3,182.0	3,177.1	2,644.0	3,729.2	0.0
ci_source	Ember	3,471.7	3,467.1	2,889.6	4,055.6	9.1
ci_source	CaDI	4,168.6	4,166.8	3,469.4	4,857.3	31.0
climate_source	ElectricityMaps	3,182.0	3,177.1	2,644.0	3,729.2	0.0
climate_source	ecoinvent	5,783.5	5,783.5	4,809.8	6,749.6	81.8

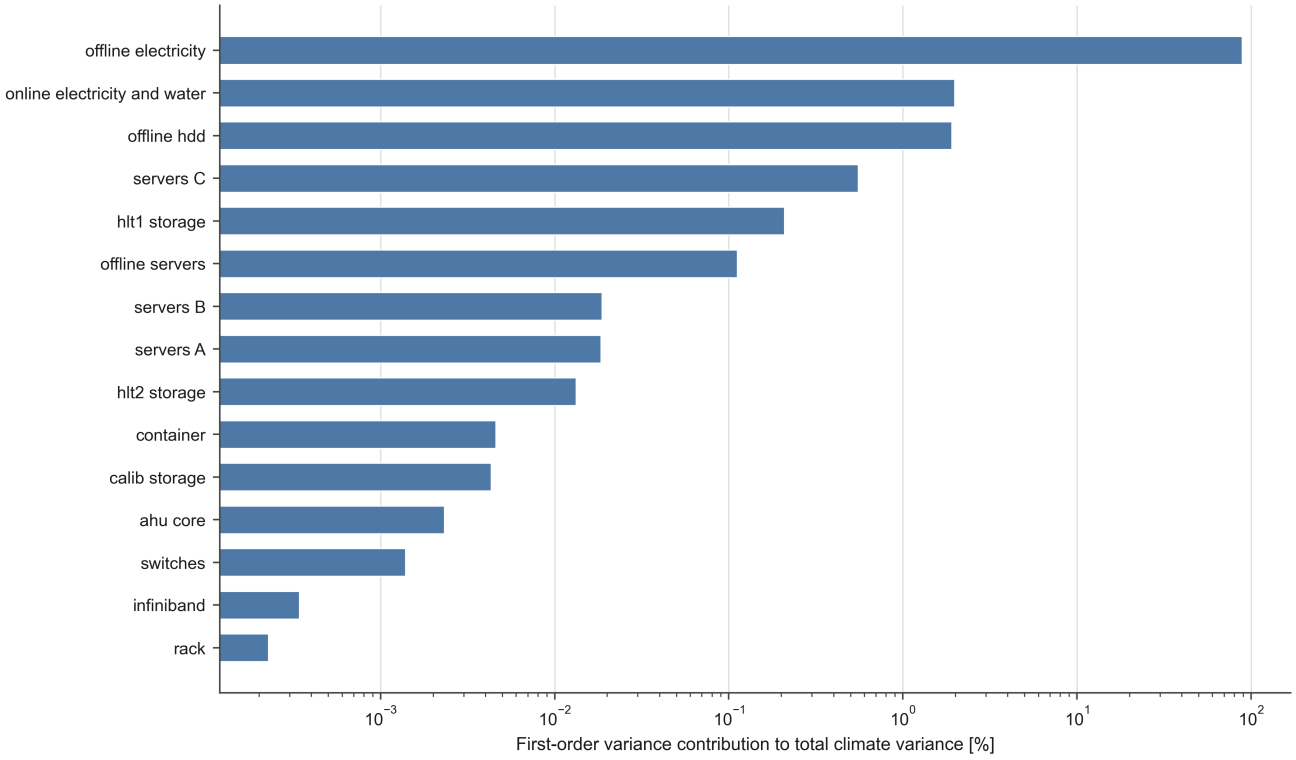


Figure 18: Leading inventory contributors to climate-result variance

Table 24: First-order climate variance contributions by component

Component	Nature	Mean [kg CO2eq]	Total variance [%]	Embodied variance [%]
offline_electricity	operational	3,142,889	88.8111	—
online_electricity_and_water	operational	1,425,084	1.9815	—
offline_hdd	embodied	456,914	1.9123	20.5098
servers_C	embodied	267,397	0.5544	5.9461
hlt1_storage	embodied	151,103	0.2090	2.2415
offline_servers	embodied	121,476	0.1121	1.2023
servers_B	embodied	49,618	0.0187	0.2006
servers_A	embodied	48,930	0.0184	0.1972
hlt2_storage	embodied	37,939	0.0132	0.1420
container	embodied	20,314	0.0046	0.0491
calib_storage	embodied	21,770	0.0043	0.0463
ahu_core	embodied	11,766	0.0023	0.0249
switches	embodied	12,348	0.0014	0.0149
infiniband	embodied	4,311	0.0003	0.0037
rack	embodied	915	0.0002	0.0024
inverter	embodied	1,659	0.0002	0.0016
water_pump	embodied	1,624	0.0001	0.0015
nvme_storage	embodied	3,836	0.0001	0.0014
transformer	embodied	976	0.0000	0.0003
fundation	embodied	1,477	0.0000	0.0002
ro_pump	embodied	271	0.0000	0.0000
ahu_air_filter	embodied	870	0.0000	0.0000
cable_signal	embodied	11	0.0000	0.0000
cable_ecs	embodied	9	0.0000	0.0000
cable_tfc	embodied	5	0.0000	0.0000

Table 25: Convergence requirements by openLCA chain

Tolerance [%]	Chain	Empirical n	Observed n	Additional n	Additional cost [h]
5	container__bg	2,138	200	1,939	9.9
5	offline_operational__bg	1,919	500	1,420	7.2
5	online_operational__bg	2,114	500	1,615	8.2
10	container__bg	535	200	335	1.7
10	offline_operational__bg	483	500	0	0.0
10	online_operational__bg	529	500	29	0.1
15	container__bg	238	200	38	0.2
15	offline_operational__bg	213	500	0	0.0
15	online_operational__bg	234	500	0	0.0
20	container__bg	130	200	0	0.0
20	offline_operational__bg	117	500	0	0.0
20	online_operational__bg	155	500	0	0.0

Table 26: Ionising-radiation convergence diagnostics

Chain	n	Median	P5	P95	P95 IC5	P95 IC95	Max / median	Max share [%]	Variance half-width [%]
container__bg	200.0	8,424	3,968	32,374	24,991	38,640	13.0	4.71	21.9
offline_operational__bg	500.0	1.52×10^{6}	834,029	5.74×10^{6}	5.08×10^{6}	6.65×10^{6}	37.3	4.96	14.5
online_operational__bg	500.0	3.35×10^{6}	1.39×10^{6}	1.29×10^{7}	1.17×10^{7}	1.70×10^{7}	17.2	2.31	11.0

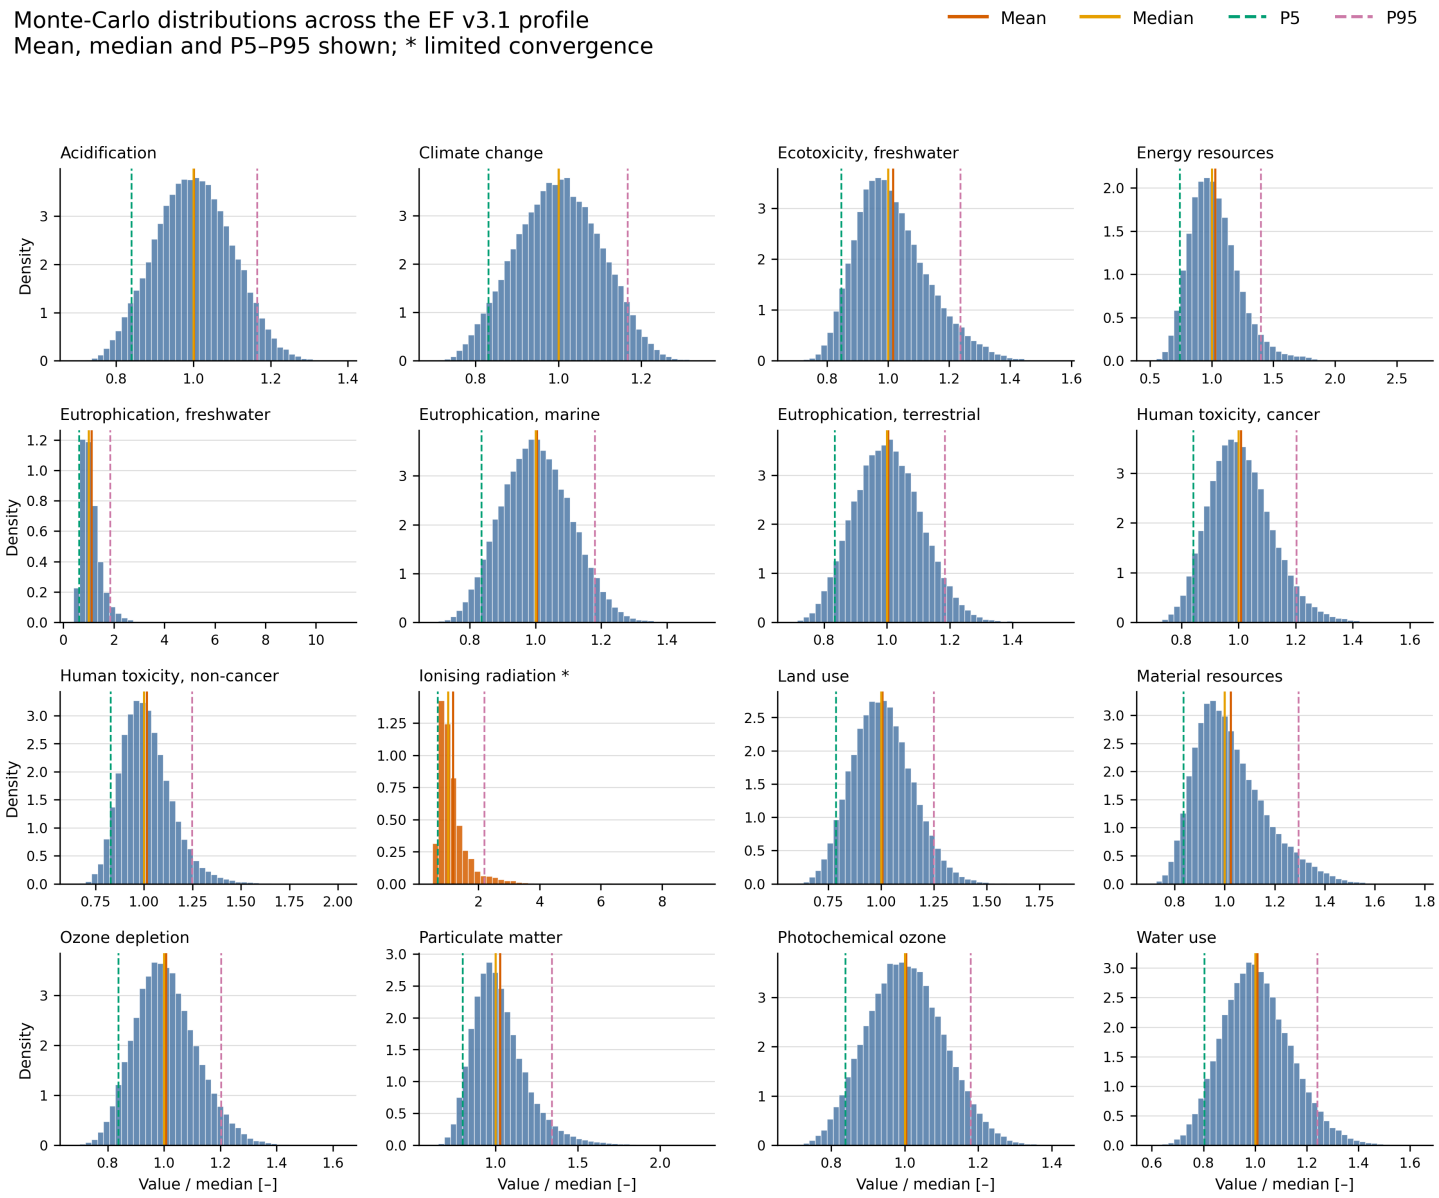


Figure 19: Complete-system Monte Carlo distributions — EF 3.1, ecoinvent 3.12 APOS

Table 27: Monte Carlo summary of annual impacts — EF 3.1

Impact category	Unit	Mean	SD	P5	Median	P95
Acidification	mol H+-Eq	24,825	2,443	20,816	24,792	28,881
Climate change	t CO_2eq	5,784	587	4,810	5,784	6,750
Ecotoxicity, freshwater	CTUe	3.41×10^{7}	3.99×10^{6}	2.84×10^{7}	3.35×10^{7}	4.15×10^{7}
Energy resources	MJ	2.27×10^{8}	4.66×10^{7}	1.65×10^{8}	2.22×10^{8}	3.10×10^{8}
Eutrophication, freshwater	kg P-Eq	2,478	1,212	1,416	2,243	4,180
Eutrophication, marine	kg N-Eq	5,758	602	4,797	5,740	6,777
Eutrophication, terrestrial	mol N-Eq	50,591	5,366	42,087	50,440	59,716
Human toxicity, cancer	CTUh	1.52×10^{-3}	1.67×10^{-4}	1.27×10^{-3}	1.51×10^{-3}	1.81×10^{-3}
Human toxicity, non-cancer	CTUh	0.056	7.23×10^{-3}	0.045	0.055	0.068
Ionising radiation	kBq U235-Eq	1.18×10^{7}	6.07×10^{6}	6.87×10^{6}	1.02×10^{7}	2.24×10^{7}
Land use	dimensionless	2.24×10^{7}	3.15×10^{6}	1.75×10^{7}	2.23×10^{7}	2.78×10^{7}
Material resources	kg Sb-Eq	108	14.8	88.0	105	136
Ozone depletion	kg CFC-11-Eq	0.187	0.021	0.156	0.186	0.224
Particulate matter	disease incidence	0.192	0.032	0.150	0.187	0.251
Photochemical ozone	kg NMVOC-Eq	15,584	1,609	13,028	15,534	18,316
Water use	m3 world eq. deprived	3.60×10^{6}	474,563	2.87×10^{6}	3.57×10^{6}	4.43×10^{6}

Results — online/offline comparison and derived indicators

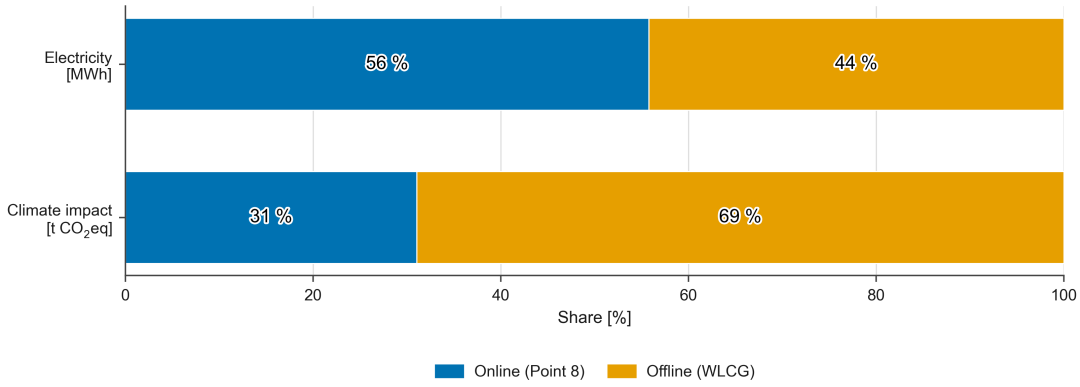


Figure 20: Online/offline electricity and climate-impact shares — Electricity Maps

Table 28: Online/offline comparison in 2025 — Electricity Maps

Quantity	Online (Point 8)	Offline (WLCG)
Installed / pledged capacity [10^9 HS23·h]	30.38	13.65
Computing work delivered [10^9 HS23·h]	—	16.25
Electricity [GWh]	12.00	9.51
Embodied impact [t CO_2eq/yr]	591	527
Operational impact [t CO_2eq/yr]	372	1,610
Total [t CO_2eq/yr]	963	2,137
Embodied share [%]	61	25
Carbon intensity, per capacity [g CO_2eq/kHS23·h]	31.7	156.5
Carbon intensity, per work delivered [g CO_2eq/kHS23·h]	—	131.5

Table 29: Climate indicators from the 2025 balance — Electricity Maps

Indicator	Central result	Unit
Total climate impact — Electricity Maps basis	3,100	t CO_2eq / year
Total climate impact — ecoinvent basis	5,734	t CO_2eq / year
Integrated luminosity recorded	11.8	fb^{-1}
Impact per unit of recorded luminosity	262.7	t CO_2eq/fb^{-1}
of which online	81.6	t CO_2eq/fb^{-1}
of which offline	181.1	t CO_2eq/fb^{-1}
Embodied share of the total	36	%

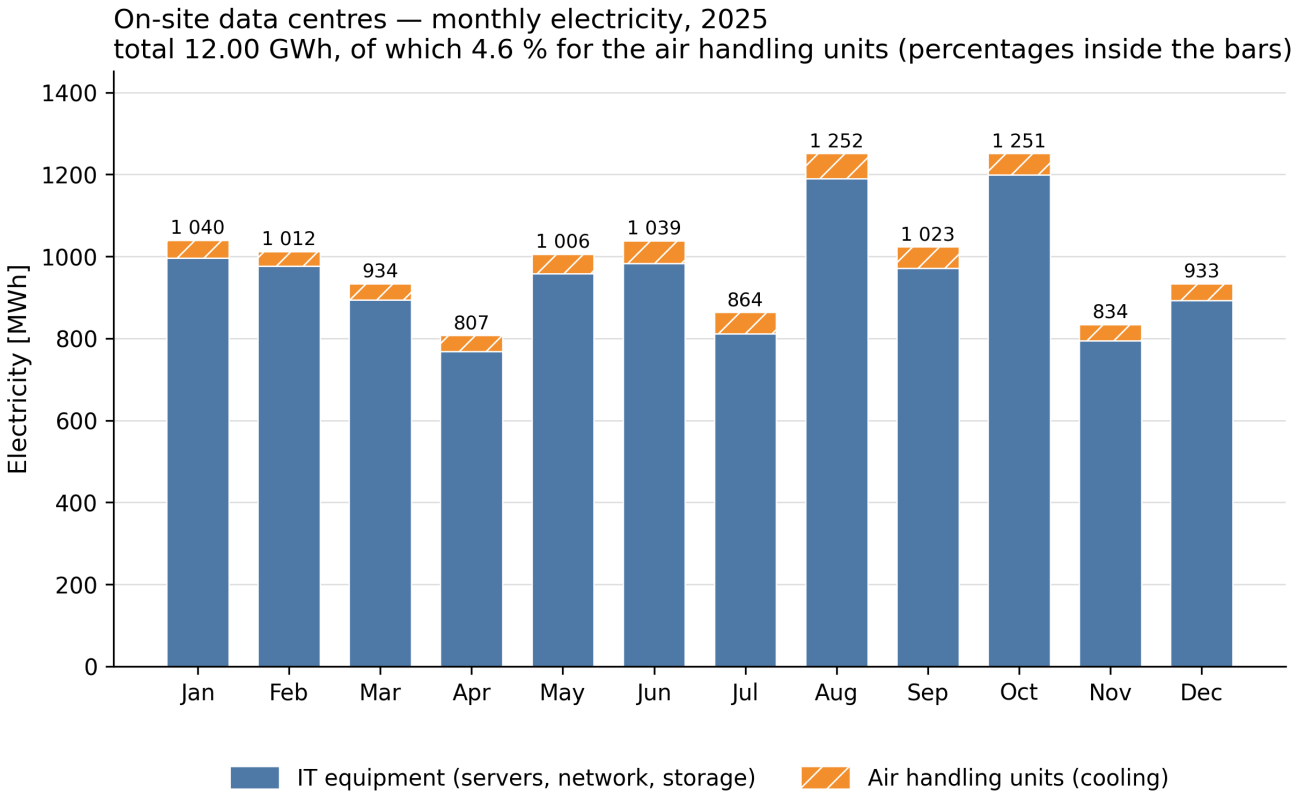


Figure 21: Monthly Point 8 electricity use in 2025

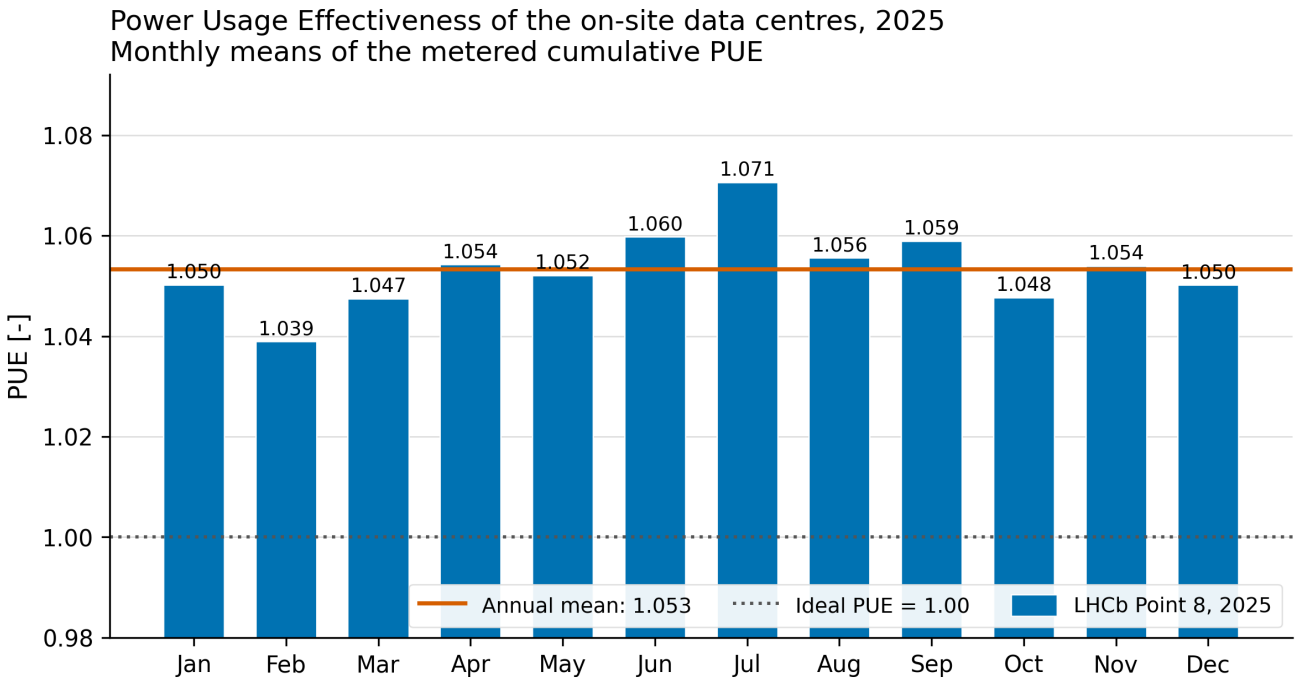


Figure 22: Monthly Point 8 PUE in 2025

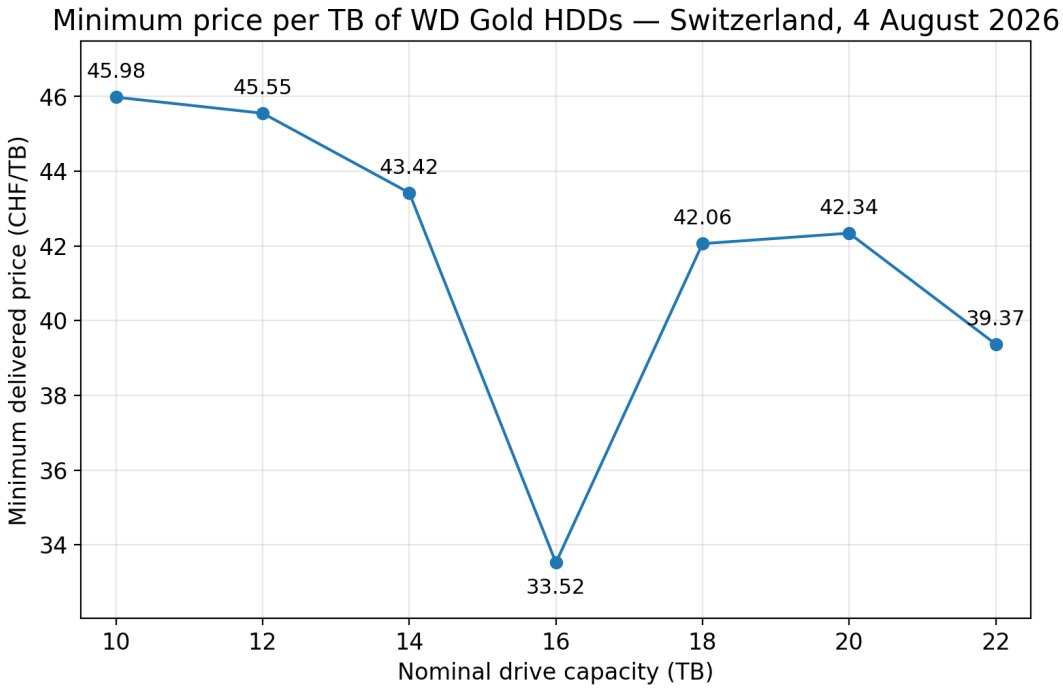


Figure 23: WD Gold delivered price per terabyte in Switzerland on 4 August 2026