

# Unveiling the Cloud’s Black Box: A Layered Model for Comparable Environmental Impact Assessment

MARIE REINBIGLER, Resilio, Switzerland

THIBAUT SIMON, Resilio, Switzerland

LOUISE AUBET, Resilio, Switzerland

At a time when public cloud usage is ever increasing due to Generative AI demand and when planetary boundaries are being exceeded more and more, it is essential to be able to understand its environmental impacts. As an answer to users’ environmental concerns, major public cloud providers offer tools for assessing the carbon emissions of cloud usage. However, their methodologies differ from one another and from existing standards. Thus, carbon footprint data from diverse public cloud providers are not comparable. Third-party solutions have emerged and offer more transparent, standards-based methodologies. However, they are limited in terms of scope, especially regarding specific cloud infrastructure features.

That is why we propose a new approach to modeling cloud services that explicitly takes into account the abstractions behind the cloud layers from Bare Metal-as-a-Service to Function-as-a-Service, as well as a more exhaustive list of infrastructure features, like load rates, virtualization overheads, or instance orchestrators. We also provide a detailed analysis of the influence of these parameters on the resulting multi-criteria environmental footprint. By demonstrating how altering these assumptions drastically changes the results, we highlight the urgent need for more standardized, transparent indicators based on physically representative data to ensure that environmental claims regarding cloud services reflect the complex reality of shared infrastructure.

CCS Concepts: • **Computing methodologies** → **Modeling and simulation**; • **Hardware** → *Impact on the environment*; • **Information systems** → **Computing platforms**.

## 1 Introduction

Human activities have undeniably driven the Earth system toward critical environmental thresholds, with the Sixth Intergovernmental Panel on Climate Change (IPCC) Report confirming our decisive role in these crises [8]. We have already transgressed seven of the nine planetary boundaries, including ocean acidification in 2025 [26]. In this context, the Information and Communication Technology (ICT) sector is a significant contributor, accounting for between 1.8% and 2.8% of greenhouse gas (GHG) emissions in 2020 [10]. In 2023, the sector consumed 27% of the carbon budget associated with the climate change planetary boundary and 18% of the boundary for mineral and metal resource use [3]. While the ICT direct environmental footprint encompasses user equipment, networks, and data centers, the latter represents a substantial portion, estimated at 20-30% of the total ICT impact [3, 10]. Furthermore, this share of impact due to data centers is increasing, mainly driven by the growing demand for cloud services and Generative Artificial Intelligence (GenAI) applications. Defined as computing resources delivered over the internet by third-party providers, public cloud services allow users to access shared resources without owning physical equipment. This paradigm encompasses various service layers, from

Infrastructure-as-a-Service (IaaS) to Software-as-a-Service (SaaS), and promotes hardware optimization and economies of scale, fueling a market projected to grow annually by 15% through 2030 [20]. However, this projected expansion comes with an important energy cost. The International Energy Agency (IEA) estimates that global data center electricity consumption will more than double, rising from approximately 415 TWh in 2024 to 950 TWh by 2030 [13], a trajectory significantly accelerated by energy-intensive applications like GenAI [9].

Despite the urgency of better mitigating these environmental impacts, assessing the footprint of cloud services remains challenging. By definition, the public cloud operates as a “black box”: users delegate and do not manage the underlying infrastructure, rendering direct physical measurement impossible and forcing reliance on models and assumptions. Furthermore, the mutualization of equipment across various clients and services requires the definition of allocation keys to distribute environmental impacts, introducing significant additional uncertainty into impact assessments.

Consequently, critical questions remain unanswered: How can the environmental claims of cloud services be trusted when modeling approaches and scopes vary massively? Which parameters drive the most significant variations in results? And fundamentally, how does the hidden underlying infrastructure dictate the environmental footprint of the cloud? This work seeks to address these questions by analyzing the methodological challenges in modeling cloud impacts, providing a new layer-based model to assess the environmental impacts of its services, and evaluating the influence of infrastructure characteristics as well as specific cloud parameters and modeling hypotheses on environmental results. The proposed model is deployed in production in Resilio Database [25].

## 2 Related works

The assessment of the environmental impacts of cloud services is currently characterized by a lack of homogeneous and consensus-based methodologies. While two primary standards exist, neither is binding, leading to a fragmented landscape where interpretation varies significantly across organizations and tools.

The first standard is ISO/IEC 21031 [17]. This specification defines the Software Carbon Intensity (SCI) score. It proposes allocation based on time and resource usage, but it lacks specific guidelines per diverse cloud services. For instance, while a virtual machine might logically be allocated based on virtual CPU (vCPU) and RAM, a database service relies more heavily on storage, thus requiring more representative allocation choices. This lack of precise prescription leaves considerable room for interpretation, which is open to very diverse estimation approaches, thus making comparison between results impossible.

Authors’ Contact Information: Marie Reinbigler, Resilio, Switzerland, marie.reinbigler@resilio-solutions; Thibault Simon, Resilio, Switzerland, thibault.simon@resilio-solutions.com; Louise Aubet, Resilio, Switzerland, louise.aubet@resilio-solutions.com.

On the other hand, the Product Category Rule (PCR) for Cloud and data center services, published by the French Environment and Energy Management Agency (ADEME), offers a more prescriptive approach for a subset of cloud services [14]. It defines specific allocation rules for typical services such as Bare Metal-as-a-Service (BMaaS), IaaS, and block storage. These rules use precise metrics including RAM capacity (GB), vCPU count, storage capacity (GB), and power consumption. While proposing a custom approach to assess non-standard cloud services by decomposing them into the defined standard components, the document formalizes formulas only up to the IaaS with limited infrastructure parameters.

Major cloud providers have developed their own carbon accounting tools. AWS provides results that are generated using the GHG protocol and Life Cycle Assessment (LCA) approach to estimate their impact, including the upstream and downstream embodied emissions [5]. Its allocation strategy is hybrid: it uses precise resource-based allocation for services relying on dedicated servers and economic allocation for services built on top of others. While the scope considered by Google is narrower, they also allocate impact to customers via physical resources consumed for major services and use monetary allocation for smaller ones [27]. Microsoft’s approach is based on monetary allocation [21]. Yet they generally do not adhere to the aforementioned standards, and their transparency levels vary widely, which makes them non-comparable.

To address the opacity of provider tools, several third-party solutions have emerged. Table 1 summarizes the scope as well as parameters and allocation approaches for two of them. *BoaviztAPI* [28] is an open-source tool targeting IaaS instances. Cloud Carbon Footprint (CCF) [31], also open-source, bases its methodology on the SCI score and relies on LCA data from *BoaviztAPI* [28] for hardware impact data. *Tailpipe* [30] is a private tool also relying on the SCI score and offers a broader scope by including network equipment and part of the technical environment. It uses *BoaviztAPI* for hardware impacts [28] as well.

A critical limitation of these works is the inadequate treatment of factors inherent to cloud infrastructure, like load rate [22], overhead [18], and compute resources orchestration. These elements, defining the efficiency and real-world resource consumption of cloud environments, are oversimplified or ignored, leading to potential underestimations in assessments. Most of them fail to capture the abstraction levels required and are therefore not suitable for modeling cloud services beyond Virtual Machines (VMs), as it requires modeling of specific additional resources.

With this work, we propose a new and more exhaustive modeling approach to assess the environmental impacts of cloud services that explicitly takes into account the abstraction behind the different cloud layers. It has already been implemented and deployed in Resilio Database [25] for users to estimate, compare, and optimize their cloud use in terms of environmental impacts.

	CCF	Boavizt-API	Tailpipe	This work
<b>Included fluxes in the scope</b>				
Computing server	✓	✓	✓	✓
GPU	✗	✗	✗	✓
Network equipment	✗	✗	✓	✗
Technical environment	(✓)	✗	(✓)	✓
Building architecture	✗	✗	✗	✓
Orchestrator	✗	✗	✗	✓
Spare equipment	✗	✗	✗	✓
<b>Parameters considered</b>				
Location	✓	✓	✓	✓
PUE	✓	✗	✓	✓
Overhead	✗	✗	✗	✓
Load rate	✗	✗	✗	✓
<b>Allocation approach</b>				
		vCPU	vCPU	vCPU
Resources used	vCPU	GB RAM	GB RAM	GB RAM
		GB sto.	GB sto.	GB sto.

Table 1. Scope, parameters, and allocation approaches for various tools. (✓) means partially taken into account.

### 3 Cloud model

#### 3.1 Life Cycle Assessment framework

LCA methodology offers a comprehensive, multi-criteria evaluation that accounts for the entire life cycle of products, services, or organizations. This method, standardized by ISO 14040 [15] and ISO 14044 [16], is largely used because of its 3 main characteristics:

- **Functional:** The object of study is defined by the function it fulfills, allowing comparison between different solutions.
- **Multi-criteria:** Several environmental indicators are taken into account. This allows to avoid impact transfer from one environmental indicator to another. In this paper, we focus on three indicators considered relevant to the ICT sector:
  - *Climate change* (Global Warming Potential (GWP)), assessing the global warming potential due to changes in the climate, measured in kg CO<sub>2</sub> equivalent [29].
  - *Resource use - minerals and metals* (Abiotic Depletion Potential for Elements (ADPe)), assessing the depletion and potential scarcity of mineral and metal resources, measured in kg antimony (Sb) equivalent [32].
  - *Water usage* (WU), assessing the impact related to freshwater consumption, measured in m<sup>3</sup> equivalent [7].
- **Life cycle approach:** The impacts across all life cycle stages of the product or service are considered, from extraction of raw materials and manufacturing, distribution, utilization, to disposal at the end of life.

#### 3.2 Modeling approach

We evaluate the environmental footprint of a cloud service based on the following functional unit: “Using a cloud service of type *X* for 1 year in country *Y*”. Our model employs a bottom-up approach to represent the stacking of cloud layers, from the physical server up to managed services (Function-as-a-Service (FaaS)). Each service layer (BMaaS, IaaS, Container-as-a-Service (CaaS), FaaS) relies on

resources from the layer below, as shown in Figure 1, and is managed by a dedicated orchestrator. Our methodology aims to determine the appropriate share of resources from each underlying layer to represent the impact attributable to a cloud service and to apply this approach layer by layer while cumulatively accounting for overhead, orchestration, and load rate effects.




Fig. 1. Structure of cloud services with virtualization layers

For a layer  $l$ , the impact of a cloud instance is decomposed into three parts:

$$I_l^{instance} = I_l^{resources} + I_l^{orchestrator} + I_l^{datacenter} \quad (1)$$

where  $I_l^{resources}$  is the impact of the requested resources within its layer on top of the underlying layers,  $I_l^{orchestrator}$  designates the impact of the orchestration of the instance layer, and  $I_l^{datacenter}$  accounts for technical environment and building impacts, which is a share of BMaaS layer one inherited by recursivity.

Because cloud infrastructures are typically not operated at full server utilization, we consider, at each layer  $l$ , the load rate  $r_l$ . In order to absorb activity peaks, cloud providers usually have  $r_l < 1$ , at the cost of idle resources. In our model, unused resources are considered to have the same energy consumption as running ones.

**Orchestrator allocation.** The orchestrator at the layer  $l$  is considered as an instance of the layer  $l - 1$  on which layer  $l$  is relying. Its impact will be shared among all the instances at layer  $l$  that are managed by this orchestrator, as displayed in Figure 2. As a consequence, we consider a compute group composed of  $n_{nodes}$  compute nodes being instances of the layer  $l - 1$ . On each of these compute nodes, there may be running  $n_{instances}$ , which gives us a total of layer  $l$  instances lying on top of the compute group to be  $n_{nodes} \times n_{instances}$ , receiving each an equal share of the orchestrator. In addition, we major the orchestrator's allocation by dividing by the load rate to take unused resources into account.

**Data center technical environment allocation.** The technical environment represents all infrastructures required for the proper functioning of computing resources. This includes all the impact of the buildings, power supply, and cooling infrastructure, except their energy consumption, which is accounted for in the definition of the BMaaS resources impact allocation with the Power Usage Effectiveness (PUE). Note that this excludes the compute hardware itself. For a given instance, we allocate its technical environment as a fraction of the total data center impact based on the instance's power consumption, using Cloud and Data Center PCR data [14].

**BMaaS resources impact allocation.** We make a distinction between BMaaS and other layers. The impact of a BMaaS instance is computed using Equation 1, with the special case of each BMaaS instance consuming the full physical server's resources. Therefore,  $I_{BMaaS}^{resources}$  is computed as the full server's impact adjusted for the load rate and the server lifespan. We consider the embodied ('emb'), i.e., the sum of manufacturing, distribution, and end-of-life, and use impacts separately:

$$I_{BMaaS}^{resources,emb} = F_{Server}^{emb} \times \frac{1 + spare\_rate}{r_{BMaaS} \times D} \quad (2)$$

with  $F_{Server}^{emb}$  the total embodied impacts of the physical server over its life cycle,  $D$  [yr] the lifespan of this server, and the spare rate, an additional fraction of equipment kept in stock in case of failure. For use impact:

$$I_{BMaaS}^{resources,use} = F_{Server}^{use} \times \frac{PUE}{r_{BMaaS} \times D} \quad (3)$$

with  $F_{Server}^{use}$  the total use impacts of the physical server during its life cycle and the PUE of the data center. The impact of the orchestrator  $I_{BMaaS}^{orchestrator}$  is computed the same way and then shared equally between all the BMaaS instances.

**Upper layer resource impact allocation.** To allocate the impact of an instance's resources within its layer, i.e.,  $I_l^{resources}$ , we consider the amount of resources it consumes among the available ones at layer  $l - 1$ , thus allocating a share  $\alpha_l^i$  of this resource impact at layer  $l - 1$ . We differentiate between the type of resources  $i$ , which can either be the amount of  $vCPU$  for compute, the  $GB$  RAM capacity for memory, the  $GB$  storage for storage or *support* equipment, meaning any other server equipment like motherboards or integrated fans.

$$\begin{aligned} I_l^{resources} &= \sum_{i \in \{vCPU, GB\ RAM, GB\ storage, support\}} I_l^{resources,i} \\ &= \sum_{i \in \{vCPU, GB\ RAM, GB\ storage, support\}} \alpha_l^i \times I_{l-1}^{resources,i} \end{aligned} \quad (4)$$

The impact allocation  $\alpha_l^i$  for the resource  $i \in \{vCPU, GB\ RAM, GB\ storage\}$  at layer  $l$  depends on the layer's *load rate*,  $r_l$ , and *overhead*,  $o_l^i$ . For each upper layer, the overhead captures the resources lost due to virtualization, containerization, and management. This reduces the compute, memory, and storage resources available to the end-user. The allocation of impacts for the resource  $i \in \{vCPU, GB\ RAM, GB\ storage\}$  at layer  $l$  is therefore defined as:

$$\alpha_l^i = \frac{w_l^i}{w_{l-1}^i \times (1 - o_l^i)} \times \frac{1}{r_l} \quad (5)$$

with  $w_l^i$  being the amount of requested resource  $i$  at layer  $l$  and  $w_{l-1}^i$  the total amount of resource  $i$  at layer  $l - 1$ .

The allocation of the *support* impacts is a weighted sum of the allocation factors from  $vCPU$ ,  $GB$  RAM, and  $GB$  storage resources defined by:

$$\alpha_l^{support} = \sum_{i \in \{vCPU, GB\ RAM, GB\ storage\}} \alpha_l^i \times \omega^i \quad (6)$$

with  $\omega^i$  values provided by the Ademe PCR Cloud [14].

As a result, the impact of underlying layers is distributed based on resource type. As one moves up the service pyramid (Figure 1),

the load rate generally increases due to improved resource sharing, while the relative overhead at each layer may vary.

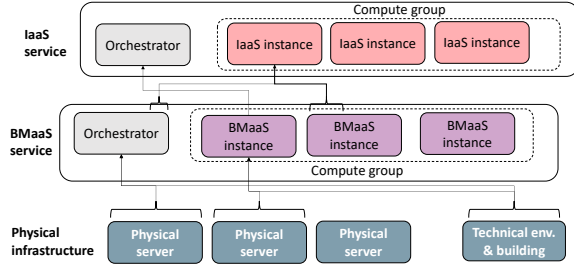


Fig. 2. Example of layer structure for an IaaS instance

This framework can be conceptually extended to handle GPUs, with some adjustments. As opposed to RAM, CPU, and storage resources, which are shared among several instances, a GPU is fully dedicated to a single instance. It thus doesn't suffer from overhead, and the allocation is performed by full units. This resource is, however, also subject to the load rate to handle peak demands of GPU resources. Equation 4 can then be extended to  $i = GPU$  with an impact allocation  $\alpha_i^{GPU}$  defined as:

$$\alpha_i^{GPU} = \frac{g_{req}}{g_{tot}} \times \frac{1}{r_i} \quad (7)$$

with  $g_{req}$  the number of GPU units requested and  $g_{tot}$  the total number of GPUs present on the server.

#### 4 Model validation

To validate our modeling approach, we compare GWP results with two tools described in Section 2: *BoaviztAPI* [6] and *Tailpipe* [30]. We compare against *Tailpipe*'s instances (AWS *c6g.metal*: 64 vCPU, 128 GiB; *c6gd.medium*: 1 vCPU, 2 GiB) [4] using a PUE of 1.04 and United Kingdom (UK) geography.

As seen in Table 2, our model yields higher impacts because it considers a broader scope encompassing load rate, overhead, spare rate, data center technical environment, and orchestration infrastructure. Consequently, we consider that for the same instance, more resources and equipment are needed. If we compare ourselves in the configuration where no load rate, overhead, spare rate, orchestrator, or data center technical environment are modeled, our methodology provides results in the same order of magnitude as the other tools. For the use impact of the *c6gd.medium* instance specifically, the

Tool	c6g.metal GWP (kg CO <sub>2</sub> eq.)		c6gd.medium GWP (kg CO <sub>2</sub> eq.)	
	Manuf.	Use	Manuf.	Use
<b>This work</b>	238.42	722.85	7.40	45.98
<b>This work (w/o orch. and infra.)</b>	139.73	491.80	2.78	14.22
<b>Tailpipe</b>	127.2	-	3.67	4.33
<b>BoaviztAPI</b>	110.58	430.0	2.13	6.70

Table 2. Comparison of GWP impacts of two IaaS instances used for 1 year with different tools.

results we obtain, without orchestrator and infrastructure, are still between 2 and 3.5 times higher than the other tools. The difference between our result and *BoaviztAPI*'s is mainly due to the absence of PUE in their computation and a difference in the estimated power consumption of the instance. The larger difference with *Tailpipe*'s result is mainly explained by a ratio of almost 2 between electricity grid mix values considered for the UK.

#### 5 Sensitivity analysis

Having validated the layer approach against existing tools at the VM level, this section explores how the approach enables identifying the most influential parameters on impact results across different cloud service levels and highlights the variability of results inherent to modeling hypotheses. In the following experiments, except if specifically mentioned, the underlying physical server is composed of an Intel Xeon Platinum 8259CL with 24 cores and 48 threads, two 256 GB RAM sticks, and one 256 GB SSD drive. The server's environmental impact has been modeled using Resilio Database<sup>1</sup>.

##### 5.1 Virtualization layer impact

Each additional abstraction layer requires infrastructure resources to manage virtualization or containerization, creating overhead. Thus, ascending the instance pyramid (Figure 1) reduces available resources. As Figure 3 illustrates, this applies to memory and storage, where fixed amounts are dedicated to overhead. For instance, with 10% overhead at the IaaS level, 3% at CaaS, and 2% at FaaS, Figure 3 shows 85.6% RAM and storage retention through the full stack: from 512 GB physical RAM to 438 GB available at the FaaS layer. The higher the overhead per layer, the lower the available resource in the upper layers.




Fig. 3. Resource availability across abstraction layers for identical underlying infrastructure, showing a reduction due to cumulative overhead at each layer.

##### 5.2 Infrastructure features impact

Figure 4 displays the impact difference when modeling infrastructure parameters (load rate, spare rate, and overhead) versus omitting them. Load rates are set to 0.7, 0.8, and 0.85 for BMaaS, IaaS, and CaaS levels, respectively [12, 19, 23]. Overhead is 10% and 3% for IaaS and CaaS levels, from industry experts, with a 5% spare rate

<sup>1</sup><https://db.resilio.tech>, version 2026.4, hash: d576f3f4a23c4ac330a96a25b449e7091e71de02409fcc44e2f75c5aa14351a1.

Component	% of manuf. impact	% of use impact
CPU	1.8	73.1
RAM	80.2	2.3
SSD	1.7	2.1
Others	16.3	22.5

Table 3. Percentage of the BMaaS impact caused by each component, decomposed between manufacturing and use phases

[14]. The impact difference increases with the abstraction level, as each layer induces additional resource loss and accounts for unused instances according to the load rate.

The difference is significant, ranging from 45.6% additional impact at the BMaaS layer to 111.1% at the CaaS layer. Consequently, incorporating these parameters with representative values is essential for accurate impact estimation of cloud services.




Fig. 4. Impact difference when modeling versus omitting infrastructure parameters (load rate, spare rate, overhead) at BMaaS, IaaS, and CaaS levels.

### 5.3 Instance resource request impact

One of the benefits of this layer-based modeling approach is that it provides granular, actionable data to reduce cloud footprint through transparent hardware allocation to cloud instances. Following the LCA methodology described in Section 3, this enables assessment of the embodied impact from hardware manufacturing against the usage impact from electricity consumption. In this section, distribution and end-of-life impacts are omitted due to their relatively low contribution compared to manufacturing and use impacts following modeling of the server's impacts in the Resilio Database.

The distribution of impact between the life cycle stages of upper layer services depends on their requested resource mix. Indeed, following Table 3, the BMaaS manufacturing impacts are dominated mainly by the RAM manufacturing impacts, while the use impacts originate primarily from CPU usage. Consequently, when requesting 2 vCPU, 4 GB RAM, 20 GB storage, depicted in Figure 6, the distribution among stages of upper layer services differs from BMaaS one, as we query only 0.7% of total server RAM capacity and 4% of vCPU capacity, thus having a lower manufacturing share than the use one. On the contrary, Figure 5 shows that when the requested vCPU, RAM, and storage are proportional to the available BMaaS

instance capacity, the distribution of impacts across stages remains consistent across all layers.

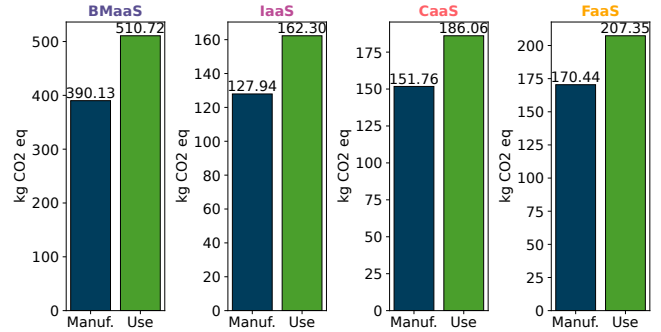


Fig. 5. Distribution of manufacturing versus use stage impacts across abstraction layers for proportional resource allocation (12 vCPU, 128 GB RAM, 64 GB storage, representing one-fourth of BMaaS capacity).

Figures 5 and 6 also highlight that, for the same resources requested at each layer, the impact increases at higher abstraction levels. Figure 7 confirms this trend for the total GWP indicator. This results from different factors. First, virtualization overhead requires larger underlying infrastructure at higher abstraction levels: a CaaS instance needs more physical resources than an equivalent IaaS instance, as additional infrastructure is dedicated to managing virtualization and containerization. Secondly, this influence is reinforced by the layer load rate, representing resources left unused to absorb activity peaks. The impacts of these unused resources are distributed among running instances. Thus, each additional layer distributes the impact of additional idle instances with higher individual impacts.

### 5.4 Underlying infrastructure impact

Throughout this paper, we have demonstrated substantial uncertainties from allocation choices of resources across cloud service layers. However, a significant source of uncertainty has been neglected: while abstraction levels can be modeled as described, the underlying physical infrastructure often remains unknown to public cloud users, further increasing assessment uncertainty.

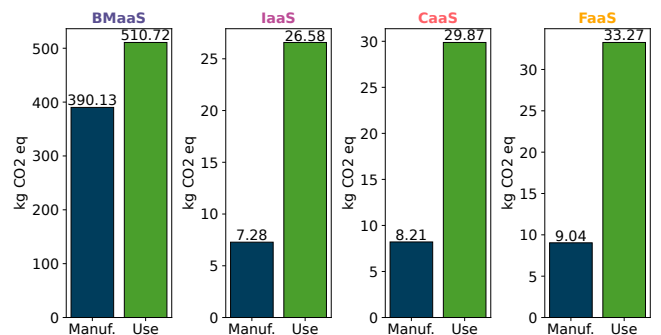


Fig. 6. Distribution of manufacturing versus use stage impacts across abstraction layers for a fixed resource request (2 vCPU, 4 GB RAM, 20 GB storage).




Fig. 7. Impact increase across abstraction layers for fixed resource allocation (2 vCPU, 4 GB RAM, 20 GB storage) on identical underlying infrastructure.

Consider two servers: a single-socket server and a dual-socket server with twice the CPU, RAM, and disk capacity (Figures 8 and 9). Despite offering twice the resources, the dual-socket server has less than double the impact at the BMaaS level for both GWP and ADPe indicators, resulting in lower impacts for equivalent IaaS, CaaS, and FaaS instances. This highlights how underlying physical infrastructure configuration significantly affects upper layer environmental impacts. Accurate physical infrastructure information is therefore essential for realistic results.




Fig. 8. Influence of underlying physical server configuration on environmental impacts across abstraction layers on the GWP indicator.




Fig. 9. Influence of underlying physical server configuration on environmental impacts across abstraction layers on the ADPe indicator.

## 6 Discussion and conclusion

In this paper, we present a new model for assessing environmental impacts of cloud services that accounts for i) the layered cloud structure up to the FaaS layer, ii) infrastructural and iii) operational components required to maintain service continuity, including layer load rate, overheads, hardware spare rate, and data center technical environment. This granular approach provides multi-criteria

results that demonstrate comparable order-of-magnitude accuracy to existing tools when evaluated at equivalent scope.

Our sensitivity analysis demonstrates that infrastructure parameters critically influence results: with default parameters, CaaS instance impacts can more than double. The underlying physical server specification affects upper layer results by up to 10%, emphasizing the need for accurate infrastructure information.

This framework is also compatible with accounting for GPU resources in the instances configuration with an allocation by units, and is in active development for a future release in Resilio Database [25].

Our layer-based approach is a step closer to a better understanding of the impact of the diverse cloud services. The transparent allocation mechanism through abstraction layers enables impact assessment of the physical infrastructure, thus revealing how manufacturing versus usage impacts shift with abstraction level and resource allocation patterns. This allows users to assess their usage impact, compare between several options, and optimize their usage by, for example, delaying workloads [2, 24]. It enables better assessing SaaS platforms' impacts that are built on top of the aforementioned cloud services and, therefore, allows their eco-design. Finally, it allows verifying cloud providers' statement that public cloud is greener than on-premise data centers [1, 11].

## References

- [1] ACCENTURE. How moving onto the aws cloud reduces carbon emissions, 2024.
- [2] ATTENNI, G., AND BARTOLINI, N. Environmentally-conscious cloud orchestration considering geo-distributed data centers, 07 2025.
- [3] AUBET, L., CHERY, S., DERREUMAUX, A., DE MONTENAY, L., PASIAN, B., AND RABOT, A. Environmental impacts of digital technology in the world, third edition, 02 2025.
- [4] AWS. Amazon ec2 c6g instances. <https://aws.amazon.com/fr/ec2/instance-types/c6g>.
- [5] AWS. Aws customer carbon footprint methodology. <https://sustainability.aboutamazon.com/aws-customer-carbon-footprint-methodology.pdf>, 10 2025.
- [6] BOAVIZTA. Multicriteria cloud instances impacts. <https://datavizta.boavizta.org/cloudimpact>, 2026. Last accessed 13 May 2026.
- [7] BOULAY, A.-M., BARE, J., BENINI, L., BERGER, M., LATHUILLIÈRE, M. J., MANZARDO, A., MARGNI, M., MOTOSHITA, M., NÚÑEZ, M., PASTOR, A. V., RIDOUTT, B., OKI, T., WORBE, S., AND PFISTER, S. The wulca consensus characterization model for water scarcity footprints: assessing impacts of water consumption based on available water remaining (aware). *The International Journal of Life Cycle Assessment* 23, 2 (Feb 2018), 368–378.
- [8] CORE WRITING TEAM, H. L., AND (EDS.), J. R. Climate change 2023: Synthesis report. contribution of working groups i, ii and iii to the sixth assessment report of the intergovernmental panel on climate change. Tech. rep., IPCC, 2023.
- [9] DE VRIES-GAO, A. The carbon and water footprints of data centers and what this could mean for artificial intelligence. *Patterns* 7, 1 (2026), 101430.
- [10] FREITAG, C., BERNERS-LEE, M., WIDDICKS, K., KNOWLES, B., BLAIR, G. S., AND FRIDAY, A. The real climate and transformative impact of ict: A critique of estimates, trends, and regulations. *Patterns* 2, 9 (2021), 100340.
- [11] GRAHAM, S., AND CLEMMONS, E. Public cloud vs enterprise datacenters — a sustainability comparison, 05 2025.
- [12] GUIDI, G., DOMINICI, F., GILMOUR, J., BUTLER, K., BELL, E., DELANEY, S., AND BARGAGLI STOFFI, F. Environmental burden of united states data centers in the artificial intelligence era, 11 2024.
- [13] IEA. Key questions on energy and ai. Tech. rep., International Energy Agency, 2026.
- [14] INRIA, A. C. Product category rule (pcr) for cloud services and data centre it hosting services, 04 2025.
- [15] ISO. Iso 14040:2006 - environmental management - life cycle assessment - principles and framework, 07 2006. Last accessed 06 May 2026.
- [16] ISO. Iso 14044:2006 - environmental management - life cycle assessment - requirements and guidelines, 07 2006. Last accessed 06 May 2026.
- [17] ISO. Iso/iec 21031:2024, 03 2024. Last accessed 29 April 2026.

- [18] JUIZ, C., AND BERMEJO, B. On the scalability of the speedup considering the overhead of consolidating virtual machines in servers for data centers. *The Journal of Supercomputing* 80, 9 (Jun 2024), 12463–12511.
- [19] KNITTEL, C. R., SENG, J. R. L., AND WANG, S. Flexible data centers and the grid: Lower costs, higher emissions?, 07 2025.
- [20] KUSHCHOV, O. Global trends in the development of cloud solutions and technologies.
- [21] MICROSOFT. A new approach for scope 3 emissions transparency. [https://download.microsoft.com/download/7/2/8/72830831-5d64-4f5c-9f51-e6e38ab1dd55/Microsoft\\_Scope\\_3\\_Emissions.pdf](https://download.microsoft.com/download/7/2/8/72830831-5d64-4f5c-9f51-e6e38ab1dd55/Microsoft_Scope_3_Emissions.pdf), 2021.
- [22] NORRIS, T. The puzzle of low data center utilization rates, 08 2025. Last accessed 09 May 2026.
- [23] PATEL, D., NISHBALL, D., AND ONTIVEROS, J. E. Ai datacenter energy dilemma - race for ai datacenter space, 03 2024. Last accessed 24 June 2026.
- [24] RADOVANOVIĆ, A., KONINGSTEIN, R., SCHNEIDER, I., CHEN, B., DUARTE, A., ROY, B., XIAO, D., HARIDASAN, M., HUNG, P., CARE, N., TALUKDAR, S., MULLEN, E., SMITH, K., COTTMAN, M., AND CIRNE, W. Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems* 38, 2 (2023), 1270–1280.
- [25] RESILIO. I want to know the environmental impact of. <https://db.resilio.tech/>.
- [26] SAKSCHEWSKI, B., CAESAR, L., ANDERSEN, L., BECHTHOLD, M., BERGFELD, L., BEUSEN, A., BILLING, M., BODIRSKY, B. L., BOTSUN, S., DENNIS, D., DONGES, J. F., DOU, X., ERIKSSON, A., FETZER, I., GERTEN, D., HÄYHÄ, T., HEBDEN, S., HECKMANN, T., HEILEMANN, A., HUISKAMP, W. N., JAHNKE, A., KAISER, J., KITZMANN, N., KRÖNKE, J., KÜHNEL, D., LAUREANTI, N. C., LI, C., LIU, Z., LORIANI, S., LUDSCHER, J., MATHESIUS, S., NORSTRÖM, A., OTTO, F., PAOLUCCI, A., POKHOTILOV, D., RAFIEZADEH SHAHI, K., RAJU, E., ROSTAMI, M., SCHAPHOFF, S., SCHMIDT, C., STEINERT, N. J., STENZEL, F., VIRKKI, V., WENDT-POTTHOFF, K., WUNDERLING, N., AND ROCKSTRÖM, J. Planetary health check 2025: A scientific assessment of the state of the planet. Tech. rep., Potsdam Institute for Climate Impact Research (PIK), Potsdam, 2025.
- [27] SCHNEIDER, I., AND MATTIA, T. Carbon accounting in the cloud: a methodology for allocating emissions across data center users, 2024.
- [28] SIMON, T., EKCHAJZER, D., BERTHELOT, A., FOURBOUL, E., RINCE, S., AND ROUYVE, R. Boaviztapi: A bottom-up model to assess the environmental impacts of cloud services. *SIGENERGY Energy Inform. Rev.* 4, 5 (Apr. 2025), 84–90.
- [29] STOCKER, T., QIN, D., PLATTNER, G.-K., TIGNOR, M., ALLEN, S., BOSCHUNG, J., NAUELS, A., XIA, Y., BEX, V., AND (EDS.), P. M. Climate change 2013: The physical science basis. contribution of working group i to the fifth assessment report of the intergovernmental panel on climate change. Tech. rep., IPCC, 2013.
- [30] TAILPIPE. Tailpipe methodology. <https://tailpipe.ai/methodology/an-overview-of-the-tailpipe-methodology/>, 2026. Last accessed 13 May 2026.
- [31] THOUGHTWORKS. Cloud carbon footprint. <https://www.cloudcarbonfootprint.org/>. Last accessed 9 May 2026.
- [32] VAN OERS, L., GUINÉE, J. B., AND HEIJUNGS, R. Abiotic resource depletion potentials (adps) for elements revisited—updating ultimate reserve estimates and introducing time series for production data. *The International Journal of Life Cycle Assessment* 25, 2 (Feb 2020), 294–308.