

# TCO<sub>2</sub>: Analyzing the Carbon Footprint of Database Server Replacements

Marc Baeuerle<sup>1</sup>, Thomas Bodner<sup>2</sup>, Martin Boissier<sup>2</sup>, Tilmann Rabl<sup>2</sup>, Ricardo Salazar-Díaz<sup>2</sup>, Florian Schmeller<sup>2</sup>, Nils Strassenburg<sup>2</sup>, Ilin Tolovski<sup>2</sup>, Marcel Weisgut<sup>2</sup>, Wang Yue<sup>2</sup>  
marcc.baeuerle@gmail.com<sup>1</sup>, firstname.lastname@hpi.de<sup>2</sup>  
Hasso Plattner Institute, University of Potsdam, Germany

## ABSTRACT

Data centers are producing a significant and increasing amount of CO<sub>2</sub> emissions. In the past, these have been predominantly due to energy generation for powering data centers. With the transition to energy sources with lower carbon production, the embodied carbon (i.e., CO<sub>2</sub> and other greenhouse gas emissions during production, transport, and end-of-life) plays an increasing role when planning server lifecycles.

While replacing an old server with newer hardware will typically reduce the power consumption of individual tasks, due to better efficiency of modern CPUs, offsetting the embodied carbon of new hardware can take months to tens of years, depending on the grid carbon intensity.

In this demo, we invite attendees to interactively analyze the ecological lifecycles of modern database servers for different workloads and grid carbon intensities. The demo enables attendees to compare servers with different CPU architectures and estimate ecological deployment cycles for database servers.

### PVLDB Reference Format:

Marc Baeuerle<sup>1</sup>, Thomas Bodner<sup>2</sup>, Martin Boissier<sup>2</sup>, Tilmann Rabl<sup>2</sup>, Ricardo Salazar-Díaz<sup>2</sup>, Florian Schmeller<sup>2</sup>, Nils Strassenburg<sup>2</sup>, Ilin Tolovski<sup>2</sup>, Marcel Weisgut<sup>2</sup>, Wang Yue<sup>2</sup>. TCO<sub>2</sub>: Analyzing the Carbon Footprint of Database Server Replacements. PVLDB, 14(1): XXX-XXX, 2020. doi:XX.XX/XXX.XX

### PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://hpides.github.io/server-buy-advisor>.

## 1 INTRODUCTION

The demand for computing resources has risen significantly over the last two decades. This increase in demand results in substantial energy consumption from both the manufacturing and operation of hardware, leading to negative environmental impacts globally.

For many years, under Dennard’s scaling principle, the CPU’s power density stayed roughly constant as transistors got smaller, enabling faster and more power-efficient CPUs [4]. Replacing hardware translated to better performance with little to no power efficiency degradation. However, after Dennard’s scaling broke down

in the mid-2000s, increasing the CPU frequency required more power usage, leading to thermal limits. CPU manufacturers have increased the core counts to keep up with the increasing performance requirements [7]. Consequently, modern CPUs incur higher manufacturing and operational carbon footprint costs than their predecessors.

This shift requires reevaluating how we plan hardware upgrades in data center environments. While upgrading to newer hardware often improves performance and energy efficiency per operation, it incurs significant upfront environmental costs [2]. Manufacturing modern CPUs, DRAM, and SSDs involves complex industrial processes, extensive raw material extraction, and high energy consumption, all of which add to the embodied carbon.

Recent work shows that emissions associated with the manufacturing process can be more than 20 times higher than those related to operating the same hardware over its lifetime [9]. Thus, even a highly efficient new server can take years to break even from an ecological point of view, i.e., offset its embodied emissions through savings in operational emissions.

This is especially important for database workloads, which are not a central focus for hardware optimizations in modern hardware. Because of the lacking focus, energy efficiency improvements are less than for standard CPU benchmark workloads. As shown in our prior work [2], replacing servers for these use cases can result in a longer break-even period, especially in regions with higher shares of low carbon electricity.

In this paper, we introduce TCO<sub>2</sub>, a tool designed to provide insights into the ecological lifecycle of a server configuration, including components such as CPU, DRAM, and storage. It evaluates these configurations against CPU and database benchmarks, taking into account specific server utilization and grid carbon intensity. The primary goal of TCO<sub>2</sub> is to assist in making informed decisions about when to upgrade hardware while prioritizing ecological efficiency.

In the demonstration, we present two use cases of TCO<sub>2</sub>. First, we calculate the total carbon footprint of a single server, based on its utilization and the local grid carbon intensity (GCI). Next, we analyze the improvements in ecological efficiency across hardware generations by comparing the servers’ performance improvement for database workloads and calculating the period necessary to offset the embodied carbon of the new hardware. TCO<sub>2</sub> is open-source and publicly available<sup>1</sup>.

The paper is organized as follows. In Section 2, we present the model to calculate the embodied and operational carbon footprint.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 14, No. 1 ISSN 2150-8097. doi:XX.XX/XXX.XX

<sup>1</sup><https://hpides.github.io/server-buy-advisor>

In Sections 3 and 4, we present the user interface of TCO<sub>2</sub> and the demonstration flow. Finally, in Section 5, we summarize the paper.

## 2 CARBON FOOTPRINT MODEL

In this demo, we determine the point in time when replacing a current server with a new server becomes environmentally beneficial, which we refer to as the ecological break-even point ( $t_{be}$ ). We first give an overview of how to find  $t_{be}$  and afterward give details on estimating a server’s power utilization.

### 2.1 Model Summary

To find  $t_{be}$ , we follow our previous work’s model [2]. A server’s carbon footprint ( $S_{CF}$ ) consist of its emissions for manufacturing ( $E_{CF}$ ) and for operating it ( $O_{CF}$ ) until the end of its life (EOL):

$$S_{CF} = E_{CF} + \sum_{t=0}^{t_{EOL}} O_{CF_t}$$

The ecological break-even point  $t_{be}$  is determined by finding the timestamp  $t$  for which the cumulative operational carbon footprint of the current server ( $c$ ) equals the sum of the new server’s ( $n$ ) embodied carbon and its cumulative operational carbon footprint:

$$\sum_{t=0}^{t_{be}} O^c_{CF_t} = E^n_{CF} + \sum_{t=0}^{t_{be}} O^n_{CF_t}$$

Since embodied carbon data from manufacturing is typically not publicly available or incomplete, we follow the framework proposed by Gupta et al. [8]. To estimate the server’s manufacturing carbon footprint ( $E_{CF}$ ) we consider the carbon footprint of its main components: CPU ( $ECF_{CPU}$ ), DRAM ( $ECF_{DRAM}$ ), and SSD ( $ECF_{SSD}$ ):

$$E_{CF} = ECF_{CPU} + ECF_{DRAM} + ECF_{SSD}$$

To estimate the operational carbon footprint ( $O_{CF}$ ), we sum the carbon emissions from operating the server’s CPUs, DRAM, and SSDs:

$$O_{CF} = OCF_{CPU} + OCF_{DRAM} + OCF_{SSD}$$

### 2.2 Power Usage and Utilization

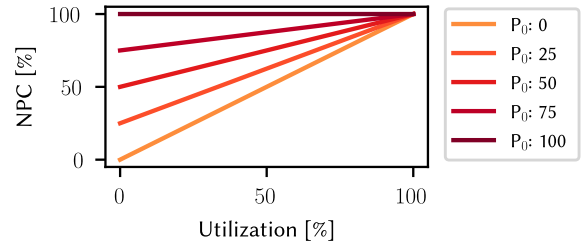
To estimate a CPU’s operational carbon emissions ( $OCF_{CPU}$ ) per year, we take the CPU’s maximum yearly power draw ( $kWh_{max/y}$ ), scale it down by a normalized power consumption (NPC) factor, reflecting the average load, and multiply the result by the grid carbon intensity (GCI) to convert power draw to carbon emissions.

$$OCF_{CPU} = kWh_{max/y} \times NPC \times GCI$$

We estimate the maximum yearly power draw ( $kWh_{max/y}$ ) by the CPU’s thermal design power (TDP), which is the power consumption under maximum theoretical load [3]. We base our normalized power consumption (NPC) on the relationship between power, energy efficiency, and utilization observed by Barroso and Hölze [1], and define a linear function based on the utilization, the CPU’s idle power draw ( $P_0$ ), and  $P_{slope}$  which represents how much the power draw increases given a 1% increase in the average yearly utilization.

$$NPC = P_0 + P_{slope} \times utilization$$

For this demo, we assume  $P_0 = 50\%$ ,  $P_{slope} = 0.5\%$ , and rely on *Electricity Maps* [6] for GCI numbers.



**Figure 1: Linear model for NPC with different parameter values for the idling power draw,  $P_0$**

Figure 1 shows the sensitivity of the NPC to the choice of values for  $P_0$ . The ideal scenario occurs when  $P_0 = 0$ ; the CPU does not draw any power when there is no work. In contrast, when  $P_0 = 100$ , the CPU draws TDP when idle.

Previous research has demonstrated that power usage for CPU-intensive workloads exhibits an exponential growth pattern as utilization increases [10]. In contrast, since database workloads are typically not CPU-intensive, a linear model is appropriate for representing the relationship between power usage and utilization.

We make the simplifying assumption in our model, that a new server will have the same utilization as a previous server and the same normalized power consumption.

## 3 TCO<sub>2</sub>: TOTAL CO<sub>2</sub> COST OF OWNERSHIP

In this paper, we introduce TCO<sub>2</sub>, a tool for quantifying the total CO<sub>2</sub> cost of ownership of database servers. We present TCO<sub>2</sub> as a web interface designed to evaluate the ecological impact of upgrading the servers through their lifecycle [2]. Currently, the tool supports the comparisons of 12 common server CPUs while allowing users to adjust other components such as DRAM, SSD, and HDD capacity along with the server workload type, utilization percentage, and electricity grid location.

### 3.1 Overview of the User Interface

Figure 2 shows the interface, which is structured into four main sections:

① **Server Configurations** – Users can configure CPU, DRAM, SSD, and HDD of an existing server to compare it with an other setup and assess its CO<sub>2</sub> emissions relative to its own embodied carbon.

② **Benchmark Settings** – This section enables users to modify the type of workload, the percentage of server utilization, and the intensity of carbon of the grid based on the location of the server.

③ **Break-Even Analysis** – The break-even time is visualized on a line chart, allowing users to assess the accumulated CO<sub>2</sub> emissions across different configurations.

④ **Detailed Breakdown** – Additional key data points, such as break-even time, grid carbon intensity, embodied carbon of new hardware, total carbon footprint until break-even, workload performance indicator, and breakdowns of the embodied and operational carbon footprint are provided to give further insights into each

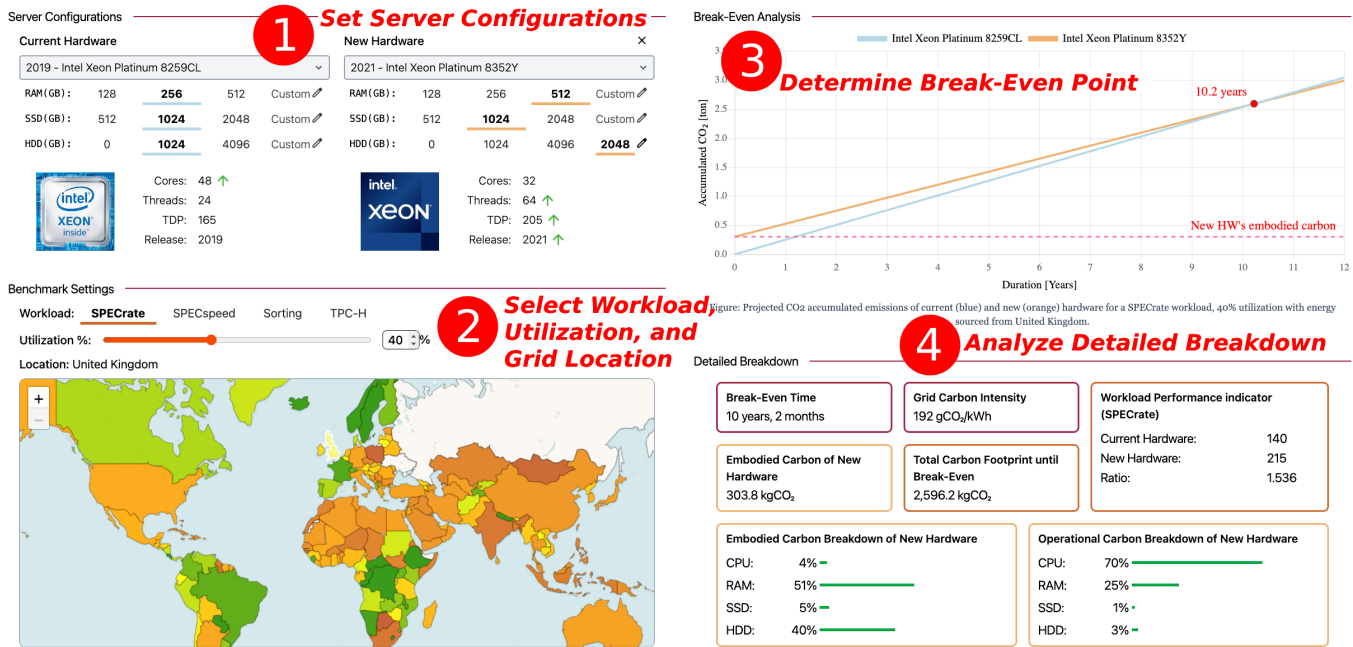


Figure 2: The user interface of TCO<sub>2</sub>.

comparison. These data points along with the line chart are then dynamically updated to reflect any changes made to the parameters.

### 3.2 Summary of Adjustable Settings

Recognizing that database workloads vary significantly, the tool provides four types of workloads to analyze CPU performance:

**SPECrate** – Measures multi-threaded performance, simulating environments such as databases and web servers [11].

**SPECspeed** – Evaluates single-threaded performance for general purpose tasks such as data compression and text processing [12]. We use publicly available measurements for both, SPECrate and SPECspeed [11, 12].

**Sorting** – A common yet computationally challenging task that is difficult to fully parallelize. A vector of four billion random integer values (`uint32_t`, 16GB) is generated, then the time to sort the entire vector is measured.

**TPC-H** – Assesses analytical database performance by running TPC-H workloads with a scale factor of 10 and 25 read-only query streams on the open-source in-memory database system Hyrise [5]. We collect measurements for Sorting and TPC-H experimentally in our lab [2].

Further settings allow users to adjust how well a server is utilized and where it is located:

**Server Utilization** – Defined as the ratio of queries per second to the maximum possible queries per second [2]. According to the findings of Barroso and Hölzle [1], who monitored thousands of Google servers over six months, servers typically operate at between 10% and 50% of their maximum theoretical capacity rather than being idle or running at peak levels [1].

**Grid Carbon Intensity (GCI)** – The GCI also plays a crucial role in predicting the ecological impact of upgrading components. The

carbon intensity of a country’s power grid measures the CO<sub>2</sub> emissions per kilowatt-hour of electricity produced. Countries with a higher share of low carbon energy, such as Sweden (22 gCO<sub>2</sub>/kWh), have significantly lower carbon intensity compared to nations like Poland (787 gCO<sub>2</sub>/kWh) [6].

### 3.3 Impact of Model Options

A key indicator of ecological efficiency improvements is the break-even time – the period required for newer hardware to provide environmental benefits when accounting for its embodied carbon and operational costs of preceding hardware [2].

The viability of upgrading a server depends on its carbon footprint and the workload performance compared to alternatives. Generally, newer CPUs outperform older ones, processing tasks more quickly and reducing energy consumption. This leads to a smaller operational carbon footprint (OCF<sub>CPU</sub>), and over time, energy savings can offset the embodied carbon produced during manufacturing, making the upgrade ecologically viable.

Our model assumes that the performance ratio between current and new CPUs reflects the efficiency gains from the upgrade. We use performance measures such as queries per second and assign performance indicators to each CPU based on workload types (SPECrate, SPECspeed, Sorting, and TPC-H) to compute this ratio:

$$\text{Ratio} = \frac{\text{New performance indicator}}{\text{Current performance indicator}}$$

Historically, new CPUs had significantly higher efficiency than previous generations and, thus, short break-even times. Newer CPUs often scale performance with additional chip space and higher power consumption, which results in longer break-even times. If

new hardware utilizes more power relative to the increase in performance, it is less efficient and, therefore, there is never a break-even point.

Server utilization and grid location heavily impact the time frame of efficient server upgrades. Increased server utilization reduces the break-even time, as more compute is used per watt, improving efficiency. In our model, the GCI is a constant multiplier on the operational emissions of the servers [2]. This results in shorter break-even points in countries with high carbon emissions as opposed to countries with less carbon-intensive energy sources.

Changes to the DRAM, SSD, and HDD capacities have an impact on the embodied and operational carbon. DRAM has the highest impact on the embodied and operational footprint.

## 4 DEMONSTRATION

The demonstration comprises two stages. First, we quantify the embodied carbon of a single server and show the time in operation necessary to offset it. We show the change in the results across different carbon intensity zones and the impact of server utilization. In the second stage, we compare the efficiency across hardware generations. We analyze the performance improvements across different workloads and quantify its impact on the servers' carbon footprint (see Figure 2). In both stages, the attendees can evaluate the hardware by configuring the CPU, memory (DRAM), and storage (SSD & HDD) capacity.

**Offsetting the embodied carbon of new hardware.** In this stage, we evaluate the embodied carbon footprint of a single server with a 2019 Intel Xeon Platinum 8259CL CPU, 256 GB of RAM, and 1TiB of SSD and HDD storage. We demonstrate the change of its carbon footprint based on the GCI in three European countries with different carbon intensities: a low GCI zone (Sweden - 22 gCO<sub>2</sub>/kWh), a medium GCI zone (UK - 192 gCO<sub>2</sub>/kWh), and a high GCI zone (Poland - 787 gCO<sub>2</sub>/kWh). The break-even point extends by 4× from Poland to the UK and by an additional 9× when moving the GCI zone from the UK to Sweden.

The attendees can observe the difference in the duration of offsetting the embodied carbon under varying server utilization (30%, 60%, 90%). We observe that the higher the utilization, the lower the duration until the embodied carbon has been offset. We observe a reduction in the break-even duration from 12-15% with each utilization increase from 30% to 60% and 90% across the three GCI zones.

**Analyzing the efficiency across hardware generations.** Each new CPU generation comes with performance and efficiency improvements advertised by the vendors. The advertised improvements are often high enough to justify the purchase of the next hardware generation. We evaluate the improvements in the context of several workloads, memory (DRAM) and storage (SSD & HDD) upgrades, the carbon intensity of the data center where the new server will operate, and the expected utilization.

The performance numbers of every new generation of CPUs have been improving for additive workloads such as SPEC CPU 2017 Speed [12] and SPEC CPU 2017 Rate [11]. We observe that the improvement factor between hardware generations decreases for sorting and database workloads (TPC-H). We demonstrate this on the 2017 Intel Xeon Platinum 8180 and 2023 Intel Xeon Platinum

8480CL. In a medium GCI zone (UK), we observe that the break-even point takes significantly longer for the database (up to 3×) and sorting (up to 7.2×) workloads.

Upgrading the memory and storage of a new server configuration has a significant impact on the new server's carbon footprint. Increasing the amount of DRAM from 256 GB to 512 GB contributes to a 2.25× increase in the duration of offsetting the embodied carbon. We observe an additional 25% increase when doubling the HDD storage capacity.

## 5 CONCLUSION

In this paper, we present TCO<sub>2</sub>: an interactive tool for calculating and analyzing the carbon footprint of database servers. We focus on determining the ecological viability of replacing existing servers with new hardware based on their embodied and operational carbon footprint. We incorporate the impact of the workload, local GCI, and the individual hardware components in our model to calculate the overall carbon footprint of a server. TCO<sub>2</sub> is open source and publicly available.

In future work, we will extend the model to cover more hardware components to cover full server deployments. Furthermore, we plan to incorporate changes in utilization for replacing a less powerful server with a larger server or consolidating deployments with placing applications running on multiple servers on fewer more powerful servers.

## REFERENCES

- [1] Luiz André Barroso and Urs Hölzle. 2007. The Case for Energy-Proportional Computing. *Computer* 40, 12 (2007), 33–37.
- [2] Thomas Bodner, Martin Boissier, Tilmann Rabl, Ricardo Salazar-Díaz, Florian Schmeller, Nils Strassenburg, Ilin Tolovski, Marcel Weisgut, and Wang Yue. 2025. A Case for Ecological Efficiency in Database Server Lifecycles. In *Conference on Innovative Data Systems Research (CIDR)*.
- [3] Intel Corporation. 2023. Thermal Design Power (TDP) in Intel Processors. Online. Retrieved 2025-03-27 from <https://www.intel.com/content/www/us/en/support/articles/000055611/processors.html>
- [4] R.H. Dennard, F.H. Gaensslen, Hwa-Nien Yu, V.L. Rideout, E. Bassous, and A.R. LeBlanc. 1974. Design of ion-implanted MOSFET's with very small physical dimensions. *IEEE Journal of Solid-State Circuits* 9, 5 (1974), 256–268. <https://doi.org/10.1109/JSSC.1974.1050511>
- [5] Markus Dreseler, Jan Kossmann, Martin Boissier, Stefan Klauk, Matthias Uflacker, and Hasso Plattner. 2019. Hyrise Re-engineered: An Extensible Database System for Research in Relational In-Memory Data Management. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*. 313–324.
- [6] Electricity Maps ApS. 2025. Live 24/7 CO<sub>2</sub> Emissions of Electricity Consumption. Online. Retrieved 2025-03-27 from <http://electricitymap.tmrow.co>
- [7] Hadi Esmailzadeh, Emily Blem, Renée St. Amant, Karthikeyan Sankaralingam, and Doug Burger. 2011. Dark silicon and the end of multicore scaling. In *2011 38th Annual International Symposium on Computer Architecture (ISCA)*. 365–376.
- [8] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu. 2022. ACT: Designing sustainable computer systems with an architectural carbon modeling tool. In *Proceedings of the Annual International Symposium on Computer Architecture (ISCA)*. 784–799.
- [9] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2022. Chasing Carbon: The Elusive Environmental Footprint of Computing. *IEEE Micro* 42, 4 (July 2022), 37–47.
- [10] Zhonghong Ou, Shichao Dong, Jiang Dong, Jukka K. Nurminen, Antti Ylä-Jääski, and Ren Wang. 2013. Characterize energy impact of concurrent network-intensive applications on mobile platforms. In *Proceedings of the International Workshop on Mobility in the Evolving Internet Architecture (MobiArch)*. 23–28.
- [11] Standard Performance Evaluation Corporation. 2025. SPEC CPU 2017 Integer Rates. Online. Retrieved 2025-03-27 from <https://www.spec.org/cpu2017/results/cpu2017/#spec-cpu-2017-integer-rates>
- [12] Standard Performance Evaluation Corporation. 2025. SPEC CPU 2017 Integer Speed. Online. Retrieved 2025-03-27 from <https://www.spec.org/cpu2017/results/cpu2017/#spec-cpu-2017-integer-speed>