# Impact of a carbon tax on the total life cycle cost of servers

Loïc Guibert (1), Olivier Weppe (2), Sébastien Rumley (3)

*(1) HES-SO University of Applied Sciences and Arts Western Switzerland, Fribourg, Switzerland. loic.guibert@hefr.ch*

*(2) Univ. Rennes, INSA Rennes, CNRS, IETR - UMR 6164, France. olivier.weppe@insa-rennes.fr*

*(3) HES-SO University of Applied Sciences and Arts Western Switzerland, Fribourg, Switzerland. sebastien.rumley@hefr.ch*

***Abstract— This paper analyzes how the ecological costs impact the total cost of ownership (TCO) of four very different servers, especially by including theoretical a carbon taxes. We assume these servers to process requests of the same type throughout their lifetime. By relating the TCO to the total number of request processes, we obtain a total cost per request (CTpR). We show that a carbon tax has a marginal effect on the CTpR when considering powerful, recent machines, whereas it can account for more than half of the costs for older or lower capacity 'frugal' equipment. However, our findings show that such 'frugal' machines are more advantageous than higher-capacity servers when considering their environmental big picture.***

***Keywords— Carbon taxes, AI, Environmental impacts, Green IT***

## I. Introduction

Information and Communication Technologies (ICT) consume significant global resources. The IEA estimated 2023 electricity demand at 1,000 TWh (~4% global total) [1], with data centers accounting for 360 TWh (1.5%). Driven by AI, data center usage is projected to double to 3% by 2030. Despite quantification challenges [2], empirical projections show a clear upward trend.

Sustainable engineering requires empirical measurement for effective carbon emissions mitigation; without granular data about effective environmental impacts, energy-saving interventions remain speculative [3]. Standard proxies like Power Usage Effectiveness (PUE) often mask the complex power profiles. Consequently, measurement is the essential first step to identify hardware inefficiencies and verify reductions [4].

As the sector's carbon footprint grows, many advocate for environmental taxes on providers [5, 6], applying the "polluter pays" principle used in industries like transportation. This study evaluates the impact of such a tax based on measured ICT service usage.

Funded by the HES-SO University of Applied Sciences and Arts Western Switzerland, for the EIAEAE project - Évaluation des Impacts Environnementaux de l'IA en Entreprise.

## II. Load Serving Power Measurements

In previous work [7], we measured the electrical energy consumption of four servers as they processed streams of identical requests with varying loads. The first server, hereafter designated as QCT, is a high-capacity AI-augmented server manufactured in 2024. This modern QCT server is compared to two more traditional servers built in 2015 (manufactured by HPE and Dell), and to a more frugal machine, namely a Raspberry Pi 4 Model B from 2021.

The testbed used to obtain measurements is displayed in Figure 1. This setup allowed us to send streams of requests with varying intensities.

Several metrics were collected while sending requests, including the CPU utilization and frequence rate, the response latency, and the power drainage through a Power Distribution Unit (PDU).

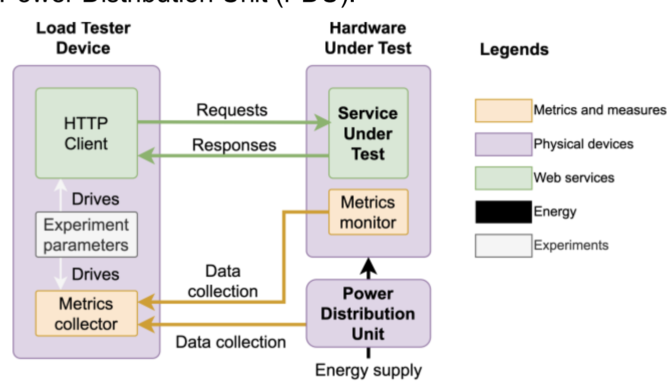


Fig 1 Testbed setup for collecting measurements

This experiment allowed us to determine, among other takeaways, that in our experimental conditions, a request to a Bert model consumes 1 to 4 Joules, if the machine is loaded close to its maximal capacity.

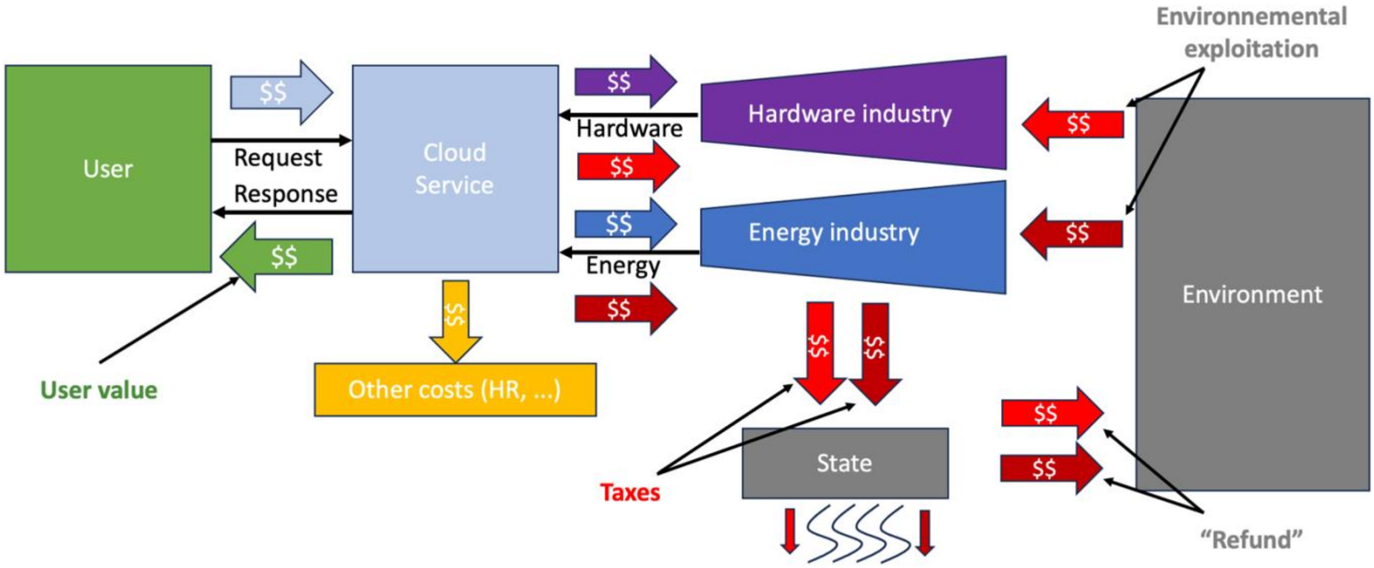


Fig 2 A model applying the polluter pays principle for the ICT industry. The users' part is displayed in green, paying for a service with an added user value. The environmental parts are displayed in gray, exploited by industries. Monetary costs are displayed in blue. The inferred environmental costs are displayed in red, appearing when a service is consumed. Those costs are propagated to the environment, where industries use this finite resource. Costs colors have shading effect to differentiate their source. The State act as a proxy that collects taxes for exploiting environmental resources, resulting in a compensation (refund) to the exploited environmental resources. Other costs are shown in yellow to complete the whole service cashflow.

## III. Server Carbon Footprint Estimation

In another piece of prior work, we studied the carbon footprint of these four servers, considering both operational emissions -related to electricity energy consumption- and embodied emissions associated with server fabrication [8]. We also analyzed how carbon footprints can be broken down into three components: their fabrication phase (embodied carbon), their static operational power when idling, and their dynamic operational power, when processing.

This work allowed us to define a novel two dimensional representation of the servers carbon emissions during their acquisition and usage phases, enabling useful insights on how to optimize their usage to lower their footprint. Observations shown that operational energy, particularly high static consumption due to limited energy proportionality in current hardware, dominates the footprint unless powered by low-carbon electricity.

## IV. Ecological Costs Model

Once measurements on server power consumption collected and their environmental footprint estimated, we can convert their carbon emissions into ecological costs.

Attributing ecological costs to parties is a complex problem, necessity rigorous rules definition and application, and extensive monitoring. The *polluter pays*

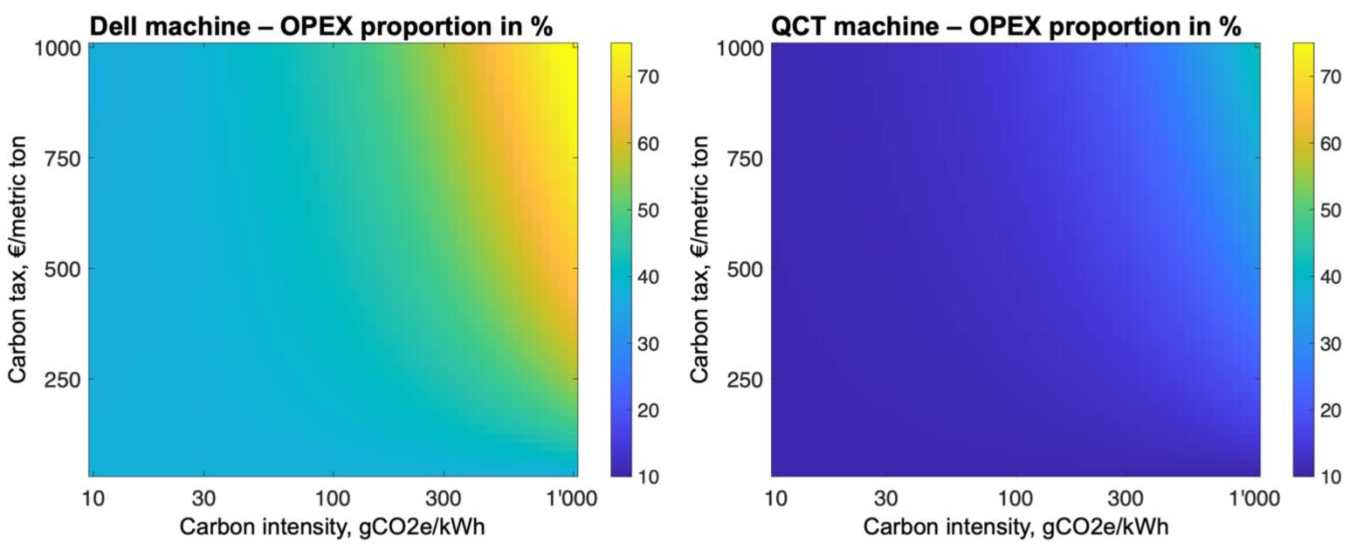


Fig 3 Distribution of operational costs at maximum occupation rates, depreciated over 8 years, and at an electricity rate of 0.20 €/kWh

*principle* is one of the most used models [9, 10], despite its known limitations, due to its accessibility.

Figure 2 shows a proposition on how to apply this principle in ICT. When users pay (blue variant arrows) to consume a Cloud service through requests and obtained responses, Cloud providers must first invest in hardware infrastructure (CAPEX) while also supplying themselves with electricity (OPEX). However, those two industries exploit environmental resources, which are finite. Cloud providers are therefore indirectly consuming environmental resources through their services. To refund environmental impacts, made through resources extraction, States can apply taxes (red variants arrows) on the industry directly exploiting environmental resources to *refund* the induced gap, offsetting non-renewable exploitation through effective actions.

## V. Introducing Carbon Taxes

On the basis of our carbon-enabled data, we now include theorical carbon taxes by applying the *polluter pays principle*, the goal being to analyze how the ecological costs impact the Total Cost of Ownership (TCO). This contribution, analyses the Total Cost per Request (CTpR), assuming that all servers are operating at full capacity and are depreciated over a certain amount of time.

Figure 3 shows the breakdown of total costs, broken down into different components, to highlight their relative weight. The machines are depreciated on 10 years. The impact of the carbon tax remains marginal for powerful, recent machines, regardless of the amount of the tax and carbon intensity. The situation is more nuanced for older or lower-capacity machines: in this case, the tax can dominate costs.

Figures 4 shows two scenarios of the costs composition per million of requests. The costs are broken down into acquisition price & acquisition taxes (CAPEX), and electricity costs & electricity taxes (OPEX). If carbon taxes are limited, the CAPEX costs always dominate, regardless of the machine capacity and acquisition price. However, a high carbon tax rate has a big impact on the costs repartition, with higher sensitivity regarding the machine maximal capacity.

As expected, the lowest-capacity machine becomes disadvantageous despite its low CAPEX, but also because its OPEX is stretched by its processing limitations. On mid-capacity machines, a high carbon tax largely influences the final costs, as the Dell and HPE machines present higher discrepancy despite similar profile and capacity.

## VI. Conclusion

The best machine choice depends highly on the considered metric. In terms of joules per request, the latest ARM-based machine is the most efficient, provided that the load is sufficient. In terms of CO2 per request, the latest machine is also the most efficient, but again, provided that demand is sufficient. In terms of real cost per request, the Dell machine gives the best results, mainly because the high price of the latest machine (QCT) cannot be fully depreciated. Our work also shows that only very high carbon taxes affect the OPEX/CAPEX ratio. The carbon taxes effect therefore highly depend on the considered machine capacities.

This entire analysis overlooks other costs that might further marginalize environmental costs.

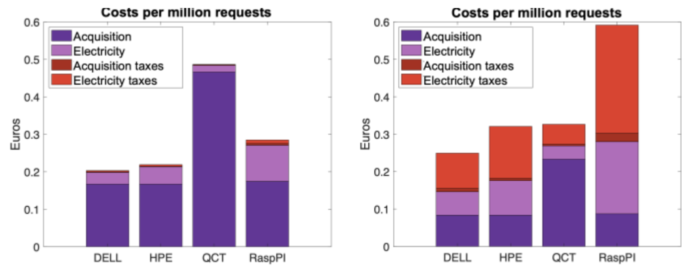


Fig 4 Costs breakdown. Scenario on the left considers 5 years depreciation, 100gCO2e/kWh, 0.1 UDC/kWh, and a carbon tax of 100 €/ton. Scenario on the right considers 10 years depreciation, 300gCO2e/kWh, 0.12 UDC/kWh, and a carbon tax of 1000 €/ton.

## References


[1] International Energy Agency. (2025). Energy and AI. https://www.iea.org/reports/energy-and-ai

[2] Bieser, J. C. T., Hintemann, R., Hilty, L. M., & Beucker, S. (2023). A review of assessments of the greenhouse gas footprint and abatement potential of information and communication technology. Environmental Impact Assessment Review, 99, 107033. https://doi.org/10.1016/j.eiar.2022.107033

[3] Alwageed, H. S., Keshta, I., Khan, R. A., Alzahrani, A., Tariq, M. U., & Ghani, A. (2024). An empirical study for mitigating sustainable cloud computing challenges using ISM-ANN. PLOS ONE, 19(11), e0308971. https://doi.org/10.1371/journal.pone.0308971

[4] Björklund, M., & Forslund, H. (2013). The purpose and focus of environmental performance measurement systems in logistics. International Journal of Productivity and Performance Management, 62(3), 230–249. https://doi.org/10.1108/17410401311309168

[5] Al-Sowaidi, A., & Isaifan, R. J. (2025). Click, Store, Emit: The Environmental Cost of Digital Infrastructure. Digital Technologies Research and Applications, 4(2), 194-211.

[6] Wang, Y., Qamruzzaman, M., & Kor, S. (2023). Greening the Future: Harnessing ICT, Innovation, Eco-Taxes, and Clean Energy for Sustainable Ecology—Insights from Dynamic Seemingly Unrelated Regression, Continuously Updated Fully Modified, and Continuously Updated Bias-Corrected Models. Sustainability, 15(23), 16417. https://doi.org/10.3390/su152316417

[7] Guibert, L., Reynaud, S., Weppe, O., Ciko, K., Welzl, M., & Rumley, S. (2025, May). TRAFFIC: Testbed foR Assessing energy eFFiciency In throughput Computing. In Proceedings of the 22nd ACM International Conference on Computing Frontiers (pp. 99-107). https://doi.org/10.1145/3706594.3727961

[8] Weppe, O., Bekri, D., Marty, T., Guibert, L., Aubet, L., Prévotet, J. C., ... & Rumley, S. (2025). Carbon Topography Representation: Improving Impacts of Data Center Lifecycle. ACM SIGENERGY Energy Informatics Review, 5(2), 41-47. https://doi.org/10.1145/3757892.3757898

[9] Luppi, B., Parisi, F., & Rajagopalan, S. (2012). The rise and fall of the polluter-pays principle in developing countries. International Review of Law and Economics, 32(1), 135-144. https://doi.org/10.1016/j.irle.2011.10.002

[10] Khan, M. R. (2015). Polluter-Pays-Principle: The Cardinal Instrument for Addressing Climate Change. Laws, 4(3), 638-653. https://doi.org/10.3390/laws4030638