Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Climate impacts of AI hardware manufacturing rival those of data centers

Sangwon Suh

sangwon@tsinghua.edu.cn

Tsinghua University

Xiaoyu Xie

Tsinghua University

Ming Xu

Tsinghua University https://orcid.org/0000-0002-7106-8390

Physical Sciences - Article

Keywords:

Posted Date: June 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9963675/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Analysis: Climate impacts of AI hardware manufacturing rival those of data centers

Xiaoyu Xie[1] *, Ming Xu*[1]*, Sangwon Suh*[1]**

[1]*School of Environment, Tsinghua University; Beijing 100084, China.*

**Corresponding author. Email: sangwon@tsinghua.edu.cn*

Data centers have been the focus of AI's electricity consumption and greenhouse gas emissions[1], while the manufacturing of AI hardware has received much less attention. Yet the advanced AI hardware that improves use-phase energy efficiency in AI data centers requires increasingly complex and energy-intensive processes and components[2]. Here, we compile a facility- and corporate-level annual greenhouse gas emissions inventory for global AI hardware manufacturing and AI data center operations. We estimate that, under market-based accounting, which accounts for renewable electricity procurement, AI hardware manufacturing emitted 5–7 Mt CO_2e yr^{-1} in 2024, exceeding the 5–6 Mt CO_2e yr^{-1} from AI data center operations. Under location-based accounting, however, operational emissions were substantially larger, at 20–21 Mt CO_2e yr^{-1} versus 6–8 Mt CO_2e yr^{-1} for AI hardware manufacturing. These findings show that understanding AI's climate footprint requires looking beyond data center electricity use to hardware manufacturing and corporate decarbonization commitments.

Introduction

The rapid expansion of artificial intelligence (AI) has raised concerns about electricity demand and associated greenhouse gas emissions[1], with data center operations becoming the most visible component of AI's climate impact[3–8]. A recent analysis based on location-specific greenhouse gas emissions intensities of nearly 10,000 data centers estimated that AI-oriented data centers

consume 53-58 TWh of electricity annually in 2024, producing 20-21 Mt CO_2e of greenhouse gas(GHG) emissions[9]. These estimates apply the greenhouse gas emission intensity of the grid regions in which AI workloads are hosted, an approach known as location-based accounting[10].

Companies can also reduce their reported Scope 2 emissions from purchased electricity through contractual instruments such as renewable electricity procurement, power-purchase agreements, electricity attribute certificates and carbon-free electricity claims[10–12]. This practice, known as market-based accounting, reflects procurement instruments in addition to the average emissions intensity of the local grid. For example, the Meta Eagle Mountain Campus facility is assigned a location-based electricity intensity of 607.2 g CO_2e kWh^{-1} in Utah, whereas Meta's disclosed Scope 2 market-based emissions factor for its data centers is 0.007 g CO_2e kWh^{-1}. This example illustrates how market-based instruments can separate procurement-adjusted reported emissions from the emissions intensity of the grids that physically serve AI workloads. Side-by-side comparison of location-based and market-based accounting is therefore useful for understanding how corporate climate actions differentiate reported emissions from the average carbon intensity of local grids.

Data center operations, however, capture only part of the AI emissions system. AI workloads depend on specialized hardware, including accelerator modules, high-bandwidth memory (HBM), advanced-node logic, high-speed interconnects, advanced packaging and server-level assembly[13–16]. The production of these components is energy- and emissions-intensive[17,18], but the collective magnitude of their global climate burden has not been clearly understood. Meanwhile, AI hardware manufacturing is becoming increasingly carbon intensive as modern AI accelerators require high-performance semiconductors, HBM[2,19–24], energy-intensive processes

such as extreme ultraviolet lithography and advanced packaging[16,22,25], and the use of potent climate-warming process-gases including fluorinated greenhouse gases for wafer fabrication and chamber-cleaning processes[26–30]. Focusing only on data center electricity use therefore understates the system-level emissions of AI and may miss potential decarbonization opportunities in AI hardware supply chains.

Here we compile a facility- and corporate-level greenhouse gas emissions inventory for AI hardware manufacturing and AI data center operations in 2024. We pair AI-oriented data center facilities with local grid factors and company-disclosed market-based electricity factors, and link AI-relevant hardware modules to semiconductor and electronics suppliers under comparable location-based and market-based accounting. This framework separates AI's physical electricity burden from procurement-adjusted reported emissions, extending AI emissions accounting from data center electricity use to major AI hardware manufacturing.

Results

Location- and market-based GHG emissions

Under location-based accounting, electricity consumption by AI data center operations emitted 19.6-21.3 Mt CO_2e yr^{-1} in 2024. We estimate that GHG emissions attributable to AI hardware manufacturing were 6.5-8.2 Mt CO_2e yr^{-1} in the same year (Fig. 1a).

Market-based accounting produces a sharply different emission hierarchy. Operational emissions decline to 5.5–6.0 Mt CO_2e yr^{-1}, while manufacturing emissions remain at 5.2–6.6 Mt CO_2e yr^{-1}, making hardware manufacturing comparable to operations and larger in the upper-bound case. This shift is driven by asymmetric reductions in disclosed emissions. Moving from location-based to market-based accounting lowers operational emissions by 14.1–15.3 Mt CO_2e yr^{-1}, or

72%, while manufacturing emissions decline by only 1.3–1.6 Mt CO_2e yr^{-1}, or 20%. This reversal reflects differences in underlying structures of GHG emissions between AI data center operators and hardware manufacturers.

For AI hardware manufacturing, whose major production hubs are located in Taiwan, China and South Korea, location-based electricity emission intensity shows a load-weighted mean of 461 g CO_2e kWh^{-1}, 29% higher than the 358 g CO_2e kWh^{-1} estimated for AI data center operations (Fig. 1b). After procurement adjustment, however, this gap widens sharply. Hardware manufacturing retains an average market-based electricity emission intensity of 363 g CO_2e kWh^{-1}, whereas operations fall to 93 g CO_2e kWh^{-1}, making the manufacturing intensity 3.9 times that of operations. (Fig. 1b). This disparity reflects a procurement gap. Major AI data center operators have more aggressively adopted onsite renewables or power purchasing agreements that use either renewable electricity directly or market-based instruments such as electricity attribute certificates (EACs) and carbon-free electricity claims, whereas AI hardware manufacturers remain concentrated in GHG-intensive production regions with lower renewable electricity coverage.

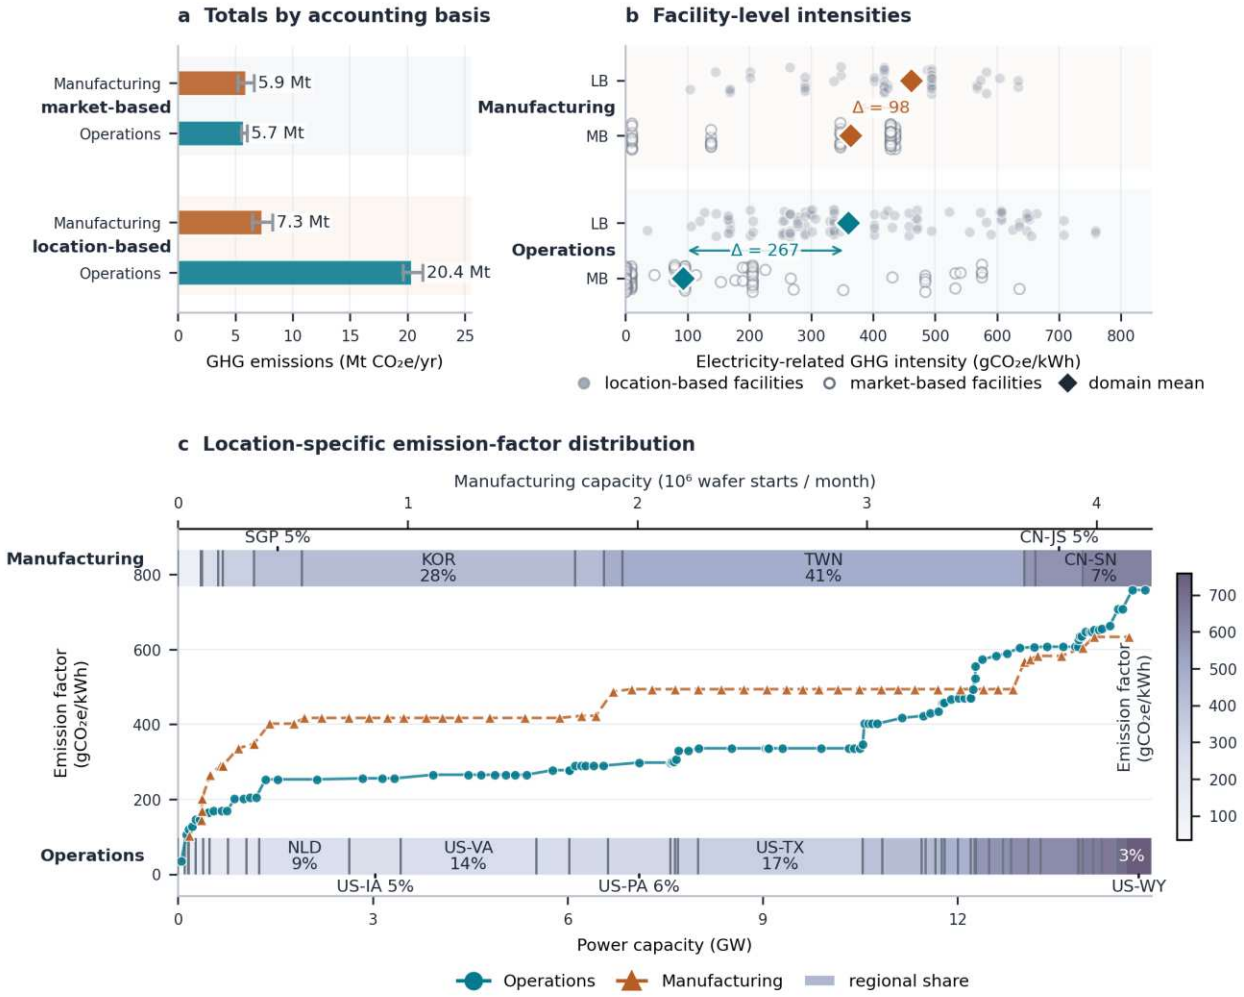

82

83 **Figure 1 | Market-based and location-based annual emissions from AI hardware**

84 **manufacturing and data center operations.**

85 **a,** Total annual GHG emissions for AI hardware manufacturing and AI data center operations under location-based and disclosed
86 market-based emission factors. **b,** Facility-level GHG intensities of electricity under location-based and disclosed market-based
87 emission factors. Points denote facilities, diamonds denote domain means and arrows indicate the reduction from location-based
88 to market-based mean intensity. **c,** Location-specific emission-factor distributions weighted by 300-mm-equivalent wafer-start
89 capacity for manufacturing and by represented AI data center power capacity for operations. Shaded bars denote regional activity
90 shares. Region abbreviations: SGP, Singapore; KOR, Korea; TWN, Taiwan; CN-JS, Jiangsu, China; CN-N, northern China;
91 NLD, Netherlands; US-IA, Iowa, United States; US-VA, Virginia, United States; US-PA, Pennsylvania, United States; US-TX,
92 Texas, United States; US-WY, Wyoming, United States.

AI hardware manufacturing is exposed to more GHG-intensive electricity than data center operations, embedding a larger share of upstream production in emission-intensive grid regions (Fig. 1c). The mapped boundary spans approximately 4.2 million 300-mm-equivalent wafer starts per month across 61 fabs or manufacturing sites and 15.0 GW of AI data center power capacity across 100 facilities or campuses, allowing the two subsystems to be compared by activity-weighted electricity exposure rather than facility counts. Under this weighting, 92.2% of manufacturing capacity lies above 400 g CO_2e kWh^{-1}, and the same share exceeds the operational load-weighted mean of 359 g CO_2e kWh^{-1}. This result reflects the main wafer-start hubs in Taiwan and Korea, which account for 41% and 28% of mapped manufacturing capacity, respectively, with additional exposure in China and Singapore. Operational capacity is more geographically dispersed, with Texas, Virginia, the Netherlands, Pennsylvania and Iowa together accounting for 7.8 GW of capacity and 32% of location-based operational emissions (Extended Data Fig. 2). These activity-weighted results show that a substantial share of AI's climate burden is embedded in geographically concentrated hardware manufacturing systems exposed to higher GHG intensity electricity.

Regional disparities in the adoption of market-based instruments

Across mapped facilities, market-based accounting reduces GHG emissions from AI data center operations by 14.7 Mt CO_2e yr^{-1}, while hardware manufacturing emissions fall by only 1.4 Mt CO_2e yr^{-1} (Fig. 2). This asymmetry reflects uneven renewable electricity procurement across supply chain (Extended Data Fig. 3).

Among the top 10% of location-based emitters among AI data centers, only 40% remain in the market-based top decile, indicating pronounced reordering of facility-level emissions rankings. The largest absolute reductions occur in North America and Europe, where market-based

operational emissions are 8.8 and 1.6 Mt CO_2e yr^{-1} lower than location-based estimates, corresponding to reductions of 86% and 68%, respectively. These reductions also align with the status of corporate-level renewable electricity, carbon-neutrality or net-zero commitments (Fig. 2). Facilities covered by explicit 2030 renewable electricity, net-zero, or carbon-neutrality commitments decline from 16.4 to 4.0 Mt CO_2e yr^{-1} under market-based accounting, a 76% reduction. Facilities lacking explicit or independently verified 2030 decarbonization pathways decline from 3.2 to 1.6 Mt CO_2e yr^{-1}, a smaller 50% reduction, and remain prominent under both accounting approaches.

By contrast, the geography of AI hardware manufacturing emissions changes little under market-based accounting. This persistence is evident in facility rankings, where 67% of facilities in the top location-based manufacturing decile remain in the corresponding market-based decile. Taiwan and Korea account for 34% and 31% of location-based manufacturing emissions, respectively, and their combined emissions decline from 4.7 to 3.7 Mt CO_2e yr^{-1} under market-based accounting. Yet they still represent 63% of market-based manufacturing emissions, indicating that supplier procurement reduces reported emissions more than it redistributes semiconductor manufacturing geography. This limited reduction in market-based accounting is consistent with 2024 supplier-level renewable electricity coverage. Major memory and foundry suppliers remain at 9–30% renewable electricity, while a high-coverage CPU supplier reaches 98% but contributes only 0.4 Mt CO_2e yr^{-1} under location-based accounting and 0.1 Mt CO_2e yr^{-1} under market-based accounting, equivalent to 5% and 2% of hardware manufacturing emissions, respectively. Market-based accounting therefore shifts the spatial prominence of AI emissions away from mature hyperscale AI data center operations, while leaving hardware manufacturing as a persistent source of AI-related emissions.

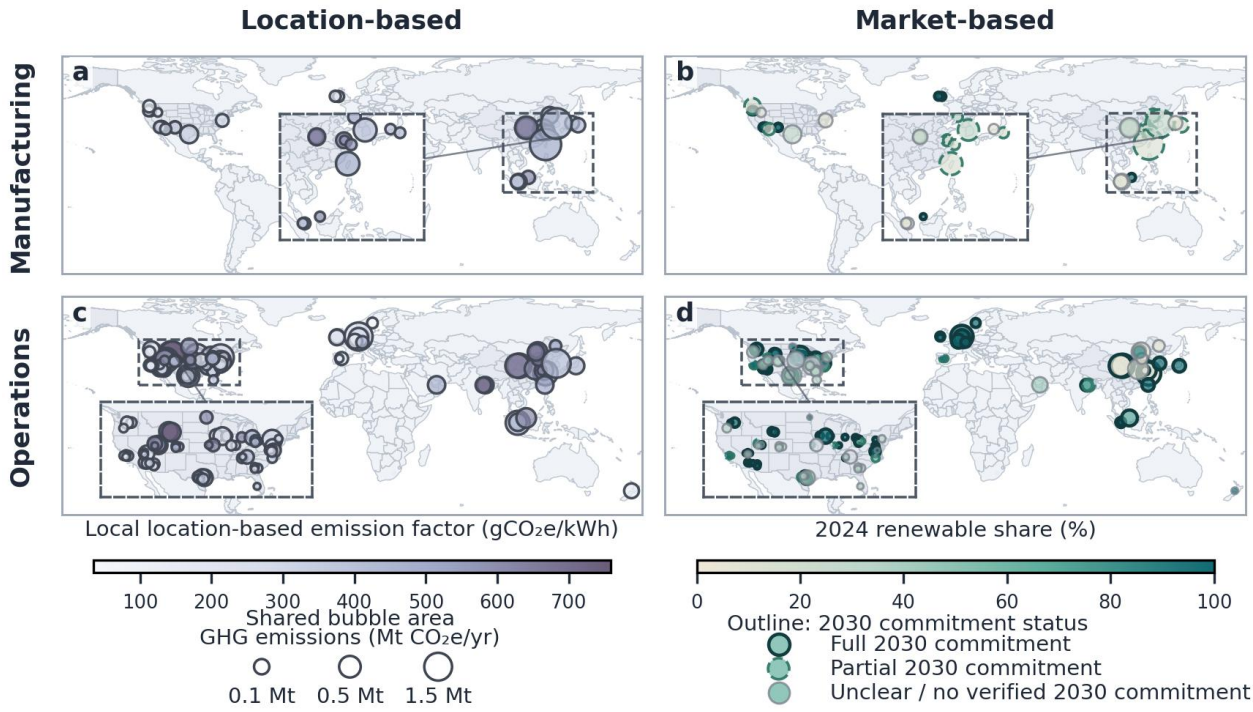

Figure 2 | Geography of location- and market-based greenhouse gas emissions of AI.

a, Location-based (LB) manufacturing emissions, with circle area denoting annual GHG emissions and circle fill denoting local LB electricity emission factors. b, Market-based (MB) manufacturing emissions, with circle fill denoting 2024 renewable electricity share and circle outline denoting 2030 commitment status. c, LB operational emissions for AI data center facilities, shown with the same bubble-area and emission-factor scales as in a. d, MB operational emissions, shown with the same renewable-share and commitment-status encodings 2030 commitment status. Full denotes a 2030-or-earlier target for 100% renewable electricity, net zero or carbon neutrality; partial denotes a time-bound target below 100%; and unclear denotes no identified 2030-or-earlier target.

Module- and corporate-level drivers

We further decompose manufacturing emissions by hardware module and operational emissions by corporate operator to identify the main contributors under each accounting basis. For hardware manufacturing, we define 12 structural module groups based on product carbon footprint data[15,31] from the Dell XE9680L and NVIDIA H100, and further refine this module

boundary according to supplier business structures and manufacturing responsibilities (Extended Data Fig. 1). Within this framework, modules 3–5 defining the accelerator-core set comprising GPU modules, HBM and NVSwitch/NVLink. Separating HBM from the broader GPU module prevents its contribution from being absorbed into aggregate accelerator categories and enables direct attribution to memory suppliers (Fig. 3a).

Accelerator-core modules—GPU modules, HBM, and NVSwitch/NVLink—dominate manufacturing emissions under both accounting bases, with HBM serving as the largest module-level contributor. HBM, GPU modules excluding HBM and NVSwitch/NVLink contribute 4.7 and 3.9 Mt CO_2e yr^{-1} under LB and MB accounting, respectively, equivalent to 64% and 66% of total manufacturing emissions. HBM alone contributes 2.4 and 1.9 Mt CO_2e yr^{-1}, accounting for 33% and 32% of total manufacturing emissions and 51% and 48% of the AI-core manufacturing burden. This concentration links the dominant module-level burden to a narrow supplier base, as HBM-attributable emissions are largely partitioned among Samsung DS, SK hynix and Micron. Samsung DS illustrates the supplier-side decarbonization gap: it contributes 1.1 Mt CO_2e yr^{-1} (59%) of market-based HBM-attributable emissions, while disclosing a 2050 net-zero target without a clearly specified 2030 clean-electricity milestone[32]. From the supplier perspective, TSMC is the largest single manufacturing-emissions contributor, with TSMC-linked GPU and NVSwitch/NVLink production contributing 1.8 Mt CO_2e yr^{-1} (31%) of market-based manufacturing emissions. Its current renewable electricity share is 14%, and its 2030 target is 60%, leaving near-term clean-electricity coverage incomplete[33]. Together, these patterns identify AI-core module concentration, HBM supplier concentration and incomplete 2030 clean-electricity coverage as the main structural determinants of manufacturing-emissions exposure.

Market-based emissions from AI data center operations show a distinct pattern of corporate procurement maturity (Extended Data Fig. 4). The top four operators in our market-based inventory—Baidu Cloud, Google, SenseTime, and China Mobile—jointly contribute 70% of market-based operational emissions, with individual contributions of 1.3, 1.2, 0.8, and 0.7 Mt CO_2e yr^{-1}, respectively. Baidu Cloud and SenseTime represent high residual contributions due to their low 2024 renewable electricity status, despite the presence of corporate-level commitments[34,35]. China Mobile shows a material residual contribution, 6% renewable electricity status and unclear 2030-or-earlier commitment coverage[36]. Colocation-dominated operators show uneven procurement maturity from their disclosures (Supplementary Material Fig. S3)[37,38], while xAI represents a behind-the-meter case[39,40].

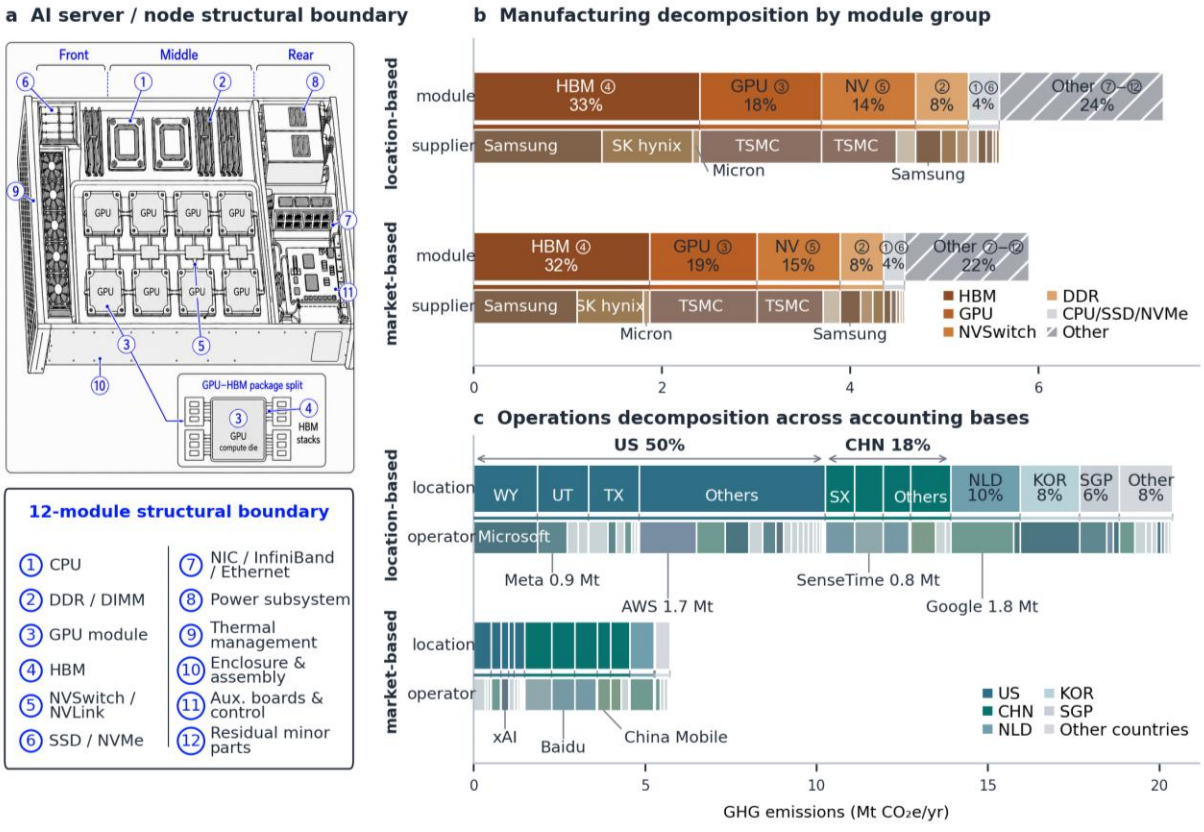

Figure 3 | Module- and operator-level GHG emissions of AI systems in 2024.

a, AI server-node boundary. Schematic of the AI server or node boundary used to define the 12 module groups included in the assessment (see in SI). b**, Manufacturing emissions by module group.** Stacked bars compare manufacturing GHG emissions across location-based and market-based accounting. Colors denote major module groups and their upstream supplier contributions. **c, Operational emissions across accounting bases.** Stacked bars compare operational GHG emissions under location-based and market-based accounting. Colors denote countries or country groups, and bars are grouped by represented AI data center power capacity.

Discussion

Our analysis shows that annual GHG emissions from AI hardware manufacturing were comparable or have exceeded AI data center operations in 2024 when facility- and corporate-level renewable electricity procurements are taken into account. Market-based manufacturing emissions are estimated at 5.9 Mt CO_2e yr^{-1}, compared with 5.7 Mt CO_2e yr^{-1} for market-based operational emissions. The location-based account preserves a different physical hierarchy, with operational emissions reaching 20.4 Mt CO_2e yr^{-1}, 2.8 times the location-based hardware-manufacturing burden. These estimates should therefore be read as a contemporaneous annual comparison of hardware-manufacturing and operational emissions, rather than as a life-cycle attribution to the same server cohort.

Under market-based accounting, reported operational emissions are substantially lower because of renewable electricity procurement, whereas hardware manufacturing remains more closely tied to GHG-intensive production grids (Extended Data Fig. 5). By linking facility-level AI infrastructure evidence with corporate procurement instruments, this study provides a dual-basis framework for distinguishing physical grid exposure from procurement-adjusted reported emissions across the AI infrastructure system.

This market-based comparison expands the mitigation agenda for AI-related GHG emissions from data center electricity use to upstream hardware supply chain. Frontier AI hardware is becoming more embodied-carbon intensive as advanced-node logic, high-bandwidth memory, larger accelerator assemblies and advanced packaging converge to deliver higher performance and greater operational efficiency. NVIDIA's disclosed product carbon footprints illustrate this trajectory. The reported cradle-to-gate footprint increases from 1,312 kg CO_2e for the HGX H100 baseboard[15] to 2,274 kg CO_2e for the B200 baseboard, a 73% increase, while the HBM share rises from 42% to 49%. This shift positions HBM as both a performance-defining component and an embodied-emissions hotspot, because memory-bandwidth scaling now carries a growing share of the manufacturing burden of frontier accelerator systems. The associated carbon-intensity risk is amplified by a concentrated supplier structure and uneven decarbonization. Samsung held the second largest HBM market share in 2024, whereas Micron remained smaller but was expanding rapidly; its HBM3E volume ramp is estimated to raise its market position toward 20–25%[41], despite a 2024 renewable electricity share of about 8%. Public commitments from Samsung and Micron remain insufficiently specific for assessing HBM-segment market-based emissions through 2030[32,42]. Without corresponding renewable electricity procurement, HBM capacity expansion could increase the market-based GHG burden of AI hardware even as data center operators continue to reduce reported operational emissions.

GHG emissions from AI data center operations decline sharply under current market-based accounting, but this reduction is highly boundary-sensitive and could be partly re-counted upward under a stricter, harmonized GHG Protocol Scope 2 framework. In our results, which are based on current disclosure records, market-based operational emissions fall by 72%. This reduction, however, depends on the methods and standards used for procurement-adjusted

electricity accounting. For example, Google, Microsoft, Meta and Amazon/AWS all report or claim 100% annual renewable electricity matching[43–46]. As AI data center electricity demand expands rapidly, sustaining these commitments may become increasingly difficult, particularly if new load growth outpaces the availability of additional, deliverable renewable supply in the relevant markets. Yet their implied market-based emissions factors differ substantially. Google has the highest factor, about 95 kg CO_2e MWh^{-1}, reflecting stricter regional matching, whereas Microsoft, Meta and Amazon/AWS report much lower factors, around 0.007–8.7 kg CO_2e/MWh under accounting boundaries that allow different combinations of power-purchase agreements, green power products, electricity attribute certificates and market-boundary hierarchies. On-going GHG Protocol revisions[10,47,48] may make market-based emissions more comparable across firms and more closely aligned with the residual carbon intensity of the grids serving corporate electricity demand.

AI emissions governance should treat data center operations and hardware manufacturing as coupled components of the same infrastructure system. A system-level disclosure framework linking facility location, electricity procurement, hardware supply chains and residual emissions would distinguish physical grid impacts from procurement-adjusted claims and identify supplier-side bottlenecks. As AI infrastructure expands, its climate implications depend not only on where models are trained and served, but also on where and how the hardware that make them possible are produced.

References

1. International Energy Agency. *Energy and AI*. (2025).

2. Schneider, I. *et al.* Life-Cycle Emissions of AI Hardware: A Cradle-To-Grave Approach and Generational Trends. *ArXiv Prepr. ArXiv250201671* https://arxiv.org/html/2502.01671v1 (2025).
3. Uptime Institute. *Global Data Center Survey 2024*. https://intelligence.uptimeinstitute.com/resource/uptime-institute-global-data-center-survey-2024 (2024).
4. Masanet, E., Lei, N. & Koomey, J. To better understand AI's growing energy use, analysts need a data revolution. https://doi.org/10.1016/j.joule.2024.07.018 doi:10.1016/j.joule.2024.07.018.
5. Masanet, E., Shehabi, A., Lei, N., Smith, S. & Koomey, J. Recalibrating global data center energy-use estimates. *Science* **367**, 984–986 (2020).
6. Alissa, H., Nick, T., Raniwala, A., & others. Using life cycle assessment to drive innovation for sustainable cool clouds. *Nature* **641**, 331–338 (2025).
7. Mytton, D. & Ashtine, M. Sources of data center energy estimates: A comprehensive review. *Joule* **6**, 2032–2056 (2022).
8. Koomey, J. G. Worldwide electricity used in data centers. *Environ. Res. Lett.* **3**, 034008 (2008).
9. Xie, X. *et al.* Climate change implications of AI depend on data center locations. *Science* (2026).
10. Greenhouse Gas Protocol. *GHG Protocol Scope 2 Guidance: An Amendment to the GHG Protocol Corporate Standard*. https://ghgprotocol.org/sites/default/files/2023-03/Scope%202%20Guidance.pdf (2015).

11. Brander, M., Gillenwater, M. & Ascui, F. Creative accounting: A critical perspective on the market-based method for reporting purchased electricity (scope 2) emissions. *Energy Policy* **112**, 29–33 (2018).
12. Xu, Q., Ricks, W., Manocha, A., Patankar, N. & Jenkins, J. D. System-level impacts of voluntary carbon-free electricity procurement strategies. *Joule* **8**, 374–400 (2024).
13. TSMC. *Climate and Nature Report*. https://esg.tsmc.com/file/public/2024-TSMC-Climate-and-Nature-Report-e.pdf (2024).
14. SK hynix. *Sustainability Report 2025*. (2025).
15. NVIDIA. *Product Carbon Footprint Summary for NVIDIA HGX H100*. https://www.nvidia.com/en-us/about-nvidia/sustainability/ (2025).
16. Kuo, C.-H., Hu, A. H., Hung, L. H., Yang, K.-T. & Wu, C.-H. Life cycle impact assessment of semiconductor packaging technologies with emphasis on ball grid array. *J. Clean. Prod.* **276**, 124301 (2020).
17. Schneider, I. & Mattia, T. Carbon accounting in the Cloud: a methodology for allocating emissions across data center users. Preprint at https://doi.org/10.48550/arXiv.2406.09645 (2024).
18. Luccioni, A. S., Viguier, S. & Ligozat, A.-L. Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. *J. Mach. Learn. Res.* **24**, 1–15 (2023).
19. Krishnan, N. *et al.* A Hybrid Life Cycle Inventory of Nano-Scale Semiconductor Manufacturing. *Environ. Sci. Technol.* **42**, 3069–3075 (2008).
20. Gupta, U. *et al.* Architectural CO2 Footprint Tool: Designing Sustainable Computer Systems With an Architectural Carbon Modeling Tool. *IEEE Micro* **43**, 107–117 (2023).

21. Williams, E. D., Ayres, R. U. & Heller, M. The 1.7 Kilogram Microchip: Energy and Material Use in the Production of Semiconductor Devices. *Environ. Sci. Technol.* **36**, 5504–5510 (2002).
22. Kuo, T.-C., Kuo, C.-Y. & Chen, L.-W. Assessing environmental impacts of nanoscale semiconductor manufacturing from the life cycle assessment perspective. *Resour. Conserv. Recycl.* **182**, 106289 (2022).
23. Boyd, S. B., Horvath, A. & Dornfeld, D. Life-Cycle Energy Demand and Global Warming Potential of Computational Logic. *Environ. Sci. Technol.* **43**, 7303–7309 (2009).
24. Schneider, I. *et al.* An Introduction to Life-Cycle Emissions of AI Hardware. *IEEE Micro* https://ieeexplore.ieee.org/iel8/40/5210076/11097303.pdf (2025).
25. Pirson, T., Delhaye, T., Pip, A. G., & others. The Environmental Footprint of IC Production: Review, Analysis, and Lessons From Historical Trends. *IEEE Trans. Semicond. Manuf.* **36**, (2023).
26. Roussilhe, G., Marquet, C. & Paris, T. From silicon shield to carbon lock-in: Electricity consumption and carbon footprint of the Taiwanese semiconductor industry. *J. Ind. Ecol.* (2023).
27. Yin, Y. & Yang, Y. Sustainable Transition of the Global Semiconductor Industry: Challenges, Strategies, and Future Directions. *Sustainability* **17**, 3160 (2025).
28. imec. *The Green Transition of the IC Industry*. https://www.imec-int.com/en/articles/green-transition-ic-industry (2022).
29. Vital, X., Krishnan, A. & Wyman, H. Pushing the Boundaries of LCA with AI Integration and Advanced Life Cycle Inventories for Semiconductor Production. in *2024 Electronics Goes Green 2024+ (EGG)* 1–4 (2024). doi:10.23919/EGG62010.2024.10631242.

30. Proske, M. & others. *LCA Critical Review Report: Life Cycle Assessment Methodology for Imec.Netzero*. (2025).
31. Dell Technologies. *Product Carbon Footprint: Generic XE9680L EU+UK*. https://www.delltechnologies.com/asset/en-us/products/servers/briefs-summaries/poweredge-xe9680l-pcf-report.pdf (2025).
32. Samsung Electronics. *Sustainability Report 2025*. (2025).
33. Taiwan Semiconductor Manufacturing Company. *Annual Report 2024*. (2025).
34. Baidu Inc. *Baidu 2024 Environmental, Social and Governance (ESG) Report*. https://esgs.cdn.bcebos.com/2025/04/28/a2/39a2aaa4f728f4e2d0aebaff346c045e.pdf (2025).
35. SenseTime Group Inc. *2024 Sustainability Report*. https://www1.hkexnews.hk/listedco/listconews/sehk/2025/0424/2025042401051.pdf (2025).
36. China Mobile Limited. *China Mobile Limited Sustainability Report 2024: Performance Report*. https://www.chinamobileltd.com/en/esg/sd/2024/08.pdf (2025).
37. Equinix. Equinix 2024 sustainability data summary. https://www.equinix.com/resources/data-sheets/sustainability-data-summary (2025).
38. Digital Realty. Digital realty published 2024 impact report, highlighting advancements in sustainability goals. https://investor.digitalrealty.com/news-releases/news-release-details/digital-realty-published-2024-impact-report-highlighting (2025).
39. Memphis Light, Gas and Water. MLGW update regarding utility support for xAI. (2025).
40. De Chant, T. xAI gets permits for 15 natural gas generators at memphis data center. https://techcrunch.com/2025/07/03/xai-gets-permits-for-15-natural-gas-generators-at-memphis-data-center/ (2025).
41. TrendForce. Micron Reportedly Targets 25% HBM Market Share by 2025.

42. Micron Technology. *Sustainability Report 2025*. https://www.micron.com/about/sustainability/sustainability-report (2025).
43. Meta Platforms, Inc. 2025 Sustainability Report. https://sustainability.atmeta.com/2025-sustainability-report/ (2025).
44. Meta. *2025 Environmental Data Index*. (2025).
45. Microsoft. *2025 Environmental Data Fact Sheet*. (2025).
46. Google. *2025 Environmental Report*. (2025).
47. Greenhouse Gas Protocol. Upcoming Scope 2 Public Consultation: Hourly Matching and Deliverability. https://ghgprotocol.org/blog/upcoming-scope-2-public-consultation-hourly-matching-and-deliverability (2025).
48. National Renewable Energy Laboratory. *Life Cycle Greenhouse Gas Emissions from Electricity Generation: Update*. https://docs.nrel.gov/docs/fy21osti/80580.pdf (2021).

Methods

Study design, inventory boundary and evidence logic

We developed a public-data greenhouse gas inventory for AI-oriented activity in 2024. The inventory contains two methodologically distinct components: AI hardware manufacturing and AI data center operations. The manufacturing component allocates entity-level greenhouse gas disclosures within corporate reporting boundaries to AI-relevant hardware activity, whereas the operations component assigns facility-level electricity use to spatially resolved location-based and market-based emission factors.

This framework extends AI emissions accounting beyond data center electricity consumption by connecting AI-relevant hardware activity to semiconductor and electronics manufacturing

entities, and by matching AI-oriented data center facilities with local grid and market-based electricity factors. The manufacturing estimate is derived from corporate greenhouse gas disclosures for five major AI-relevant manufacturing entities, TSMC, Intel, Samsung Electronics DS, Micron, and SK hynix. The mapped manufacturing coverage spans approximately 4.2 million 300-mm-equivalent wafer starts per month across 61 fabs or manufacturing sites, and the mapped operational coverage spans 15.0 GW of AI data center power capacity across 100 facilities or campuses. These activity measures provide a shared scale reference for the two analytical components. Full component definitions, evidence codes, and source records are provided in the Supplementary Information and Supplementary Data.

Dual accounting basis and market-based evidence hierarchy

All manufacturing and operational results are reported using both location-based and market-based accounting. The location-based basis represents the physical geography of production or computing and the associated grid or local energy-supply evidence. The market-based basis represents procurement-adjusted electricity accounting, including disclosed contractual electricity instruments, renewable electricity matching and reported Scope 2 market-based information where available. We therefore interpret market-based results as disclosure-based reported emissions, not as a direct measure of physical avoided emissions.

We maintain comparability by holding the underlying AI activity boundary fixed across accounting bases. For operations, the same facility electricity values are multiplied by either location-based or market-based emission factors. For manufacturing, the same hardware boundary, supplier mapping and facility or site evidence are retained, while corporate or site-linked emissions data are assigned on either a location-based or market-based basis.

Location-based estimates provide the grid-location counterpoint needed to evaluate the physical electricity burden of AI activity. Market-based data are assigned using a disclosed-first evidence hierarchy common to both manufacturing and operations. The hierarchy governs how market-based information is assigned to facility rows, manufacturing-site rows, company-level records and supplier-linked activity records.

Tier 1 uses directly disclosed data center or operations market-based emission factors. Tier 2 uses company-level market-based electricity factors when they can be attributed to the operator's data center electricity boundary. Tier 3 uses renewable-share-derived factors only when no disclosed market-based factor is available, but an attributable renewable electricity share can be linked to the relevant facility, site, company or activity boundary. For electricity-factor applications, the assigned market-based factor is calculated as

$$\phi_f^{MB} = \phi_f^{LB}\left(1 - R_f\right),$$

where f indexes the relevant facility or site record, ϕ_f^{LB}s the location-based electricity emission factor and R_f is the attributable renewable electricity share.

Tier 4 covers entities with no clear attributable market-based electricity emission factor. These entities may disclose mitigation actions, local-supply evidence or on-site generation evidence, while lacking a quantified factor assignable to the relevant electricity boundary. For these cases, the location-based factor is retained as the market-based factor:

$$\phi_f^{MB} = \phi_f^{LB}.$$

This treatment follows the logic of the GHG Protocol Scope 2 market-based method, in which electricity emissions may reflect contractual instruments, supplier-specific electricity products, renewable electricity claims and other disclosed procurement mechanisms. The resulting market-

based factors are interpreted as procurement-adjusted electricity emission factors under a recognized Scope 2 accounting framework. They are reported alongside location-based factors so that contractual electricity accounting and grid-location electricity burden can be evaluated together.

AI data center operations

The operations module estimates AI electricity use and associated emissions for operational AI data centers in 2024 that are linked to screened AI operators. Location-based electricity use and emissions are directly adopted from our previous facility-level modelling framework[9], in which low, central and high estimates were generated using a parameter-side Monte Carlo approach. That framework held route structure fixed, retained directly observed power routes as fixed inputs, and propagated uncertainty in imputation-dependent routes by sampling parameters such as realization fractions, commissioning years, ramp duration, load factors, imputed power-usage effectiveness and power terms from stratified donor pools or bounded parametric distributions.

Let f index an AI data center facility record and let $c(f)$ denote the company matched to that record. For each facility row, E_f is annual AI-attributable electricity, ϕ_f^{LB} is the location-based emission factor assigned from the facility location or resolved local energy supply, and ϕ_f^{MB} is the market-based emission factor assigned through the tiered evidence protocol described below. Operational emissions are calculated by applying the corresponding factor to the same facility electricity value:

$$G_{\text{ops}}^{LB} = \sum_f E_f \, \phi_f^{LB}.$$

$$G_{\text{ops}}^{MB} = \sum_f E_f \, \phi_f^{MB}.$$

The total operational electricity boundary and the effective operational emissions intensity are calculated after aggregation:

$$E_{\text{ops}} = \sum_{f} E_f.$$

$$I^{b}_{\text{ops}} = \frac{G^{b}_{\text{ops}}}{E_{\text{ops}}}, \qquad b \in \{\text{LB,MB}\}.$$

The market-based operational estimate therefore differs from the location-based estimate through the assigned market-based factor. The facility electricity boundary, company screening logic and AI-attributable electricity values remain unchanged.

AI hardware manufacturing

The manufacturing module estimates AI-attributable emissions from hardware manufacturing activity linked to AI server and accelerator systems. The module is defined at the AI server or node level using a 12-module structural decomposition. The full boundary includes CPU, DDR/DIMM, GPU module excluding HBM, HBM, NVSwitch/NVLink, SSD/NVMe, NIC/InfiniBand/Ethernet, power subsystem, thermal management, mechanical enclosure and assembly, auxiliary boards and control electronics, and residual minor parts.

The main assigned module groups are host CPU, accelerator logic excluding HBM, HBM, system memory represented by DDR/DIMM, and local storage represented by SSD/NVMe. NVSwitch/NVLink is treated as a logic-die extension.

Let i index selected manufacturing entities, m index AI hardware modules, b denote the accounting basis and k denote the low, central or high case. For entity i, $C_{i,b}$ denotes the reported 2024 corporate emissions envelope under accounting basis b. The allocation model combines a

module share $w_{m,k}$, a supplier-evidence mapping share $s_{i,m,k}$ and a coverage modifier $q_{i,m,k}$. The AI-attributable allocation shares for entity i is

$$\alpha_{i,k} = \sum_m w_{m,k} \times s_{i,m,k} \times q_{i,m,k}$$

The manufacturing estimate for accounting basis b and case k is then

$$\mathrm{GHG}_{\mathrm{mfg}}^{b,k} = \sum_i C_{i,b} \times \alpha_{i,k}$$

The same module structure, supplier-evidence mapping and coverage logic are used for location-based and market-based manufacturing estimates. Differences between the two estimates arise from the corporate emissions envelope used for accounting basis b, including disclosed market-based Scope 2 reporting. Market-based electricity claims are applied only within the reporting boundary of the entity that disclosed them.

For manufacturing, low, central and high cases are allocation-side scenario settings, not re-estimates of the corporate emissions envelope. For a given accounting basis, case k changes the AI-relevant allocation parameters, including module shares, supplier-evidence mapping, coverage modifiers and unallocated hardware treatment.

Acknowledgements

Acknowledgments: During preparation of this work, we used AI-assisted tools to support grammar and language editing and to help standardize and clean portions of the underlying dataset. All cleaned data, analytical outputs, and manuscript text were subsequently reviewed and validated by the authors, who take full responsibility for the accuracy and integrity of the work.

Funding: SS was supported by Xinghua Chair funds.

Author contributions:

Conceptualization: SS

Methodology: SS, XX

Investigation: XX

Formal analysis: XX

Data curation: XX

Supervision: MX, SS

Visualization: XX

Writing – original draft: XX, SS

Writing – review & editing: MX, SS

Competing interests: Authors declare that they have no competing interests.

Data availability: All data supporting the findings of this study are provided in the Article, Extended Data, Supplementary Information and Supplementary Data. The facility-level operations inventory, location-based emission-factor assignments, market-based factor assignments, manufacturing module-to-supplier mappings, scenario outputs and figure source data are provided in Supplementary Data 1. The public source documents and datasets used to construct the inventory are listed in Supplementary Data 1 and cited in the Article or Supplementary Information. No restricted-access or proprietary datasets were used.

Code availability: The custom code used for data harmonization, emissions calculations, uncertainty analysis and figure generation is provided for review in the submitted code archive, AI-Hardware-GHG-Accounting_review_package.zip. The archive includes the release version of the repository, processing scripts, figure-generation scripts and documentation needed to

reproduce the reported results. A public archived version with a persistent identifier will be made available upon publication.

Materials availability: No unique physical materials were generated or used in this study.

Extended Data legends

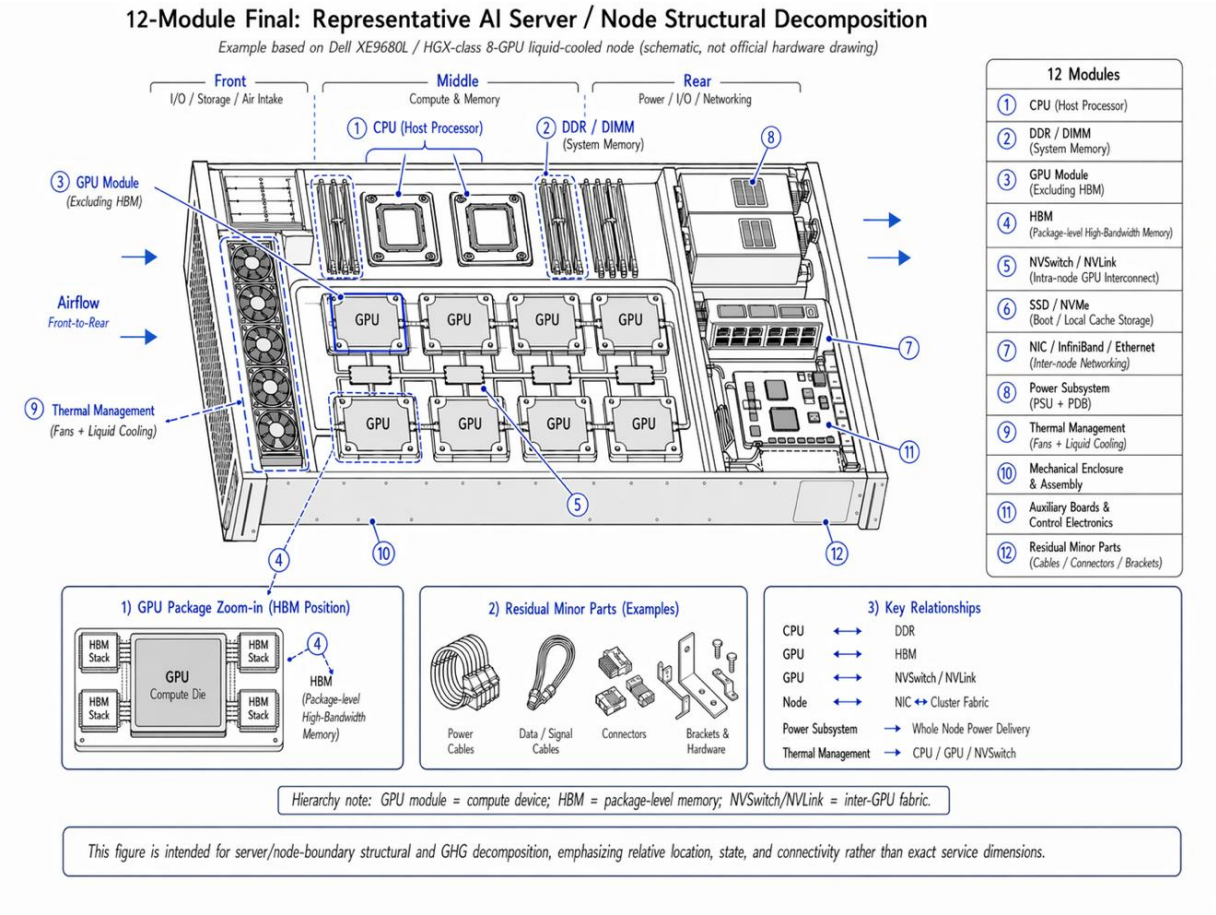

Extended Data Fig. 1 | AI server/node hardware boundary. *Schematic of the 12-module server/node boundary used for allocation and diagnostics. Modules are assigned to selected manufacturing disclosure entities only when public production-route, supplier-share or process evidence supports the mapping; otherwise, they remain in the boundary as an unallocated, boundary-preserving remainder.*

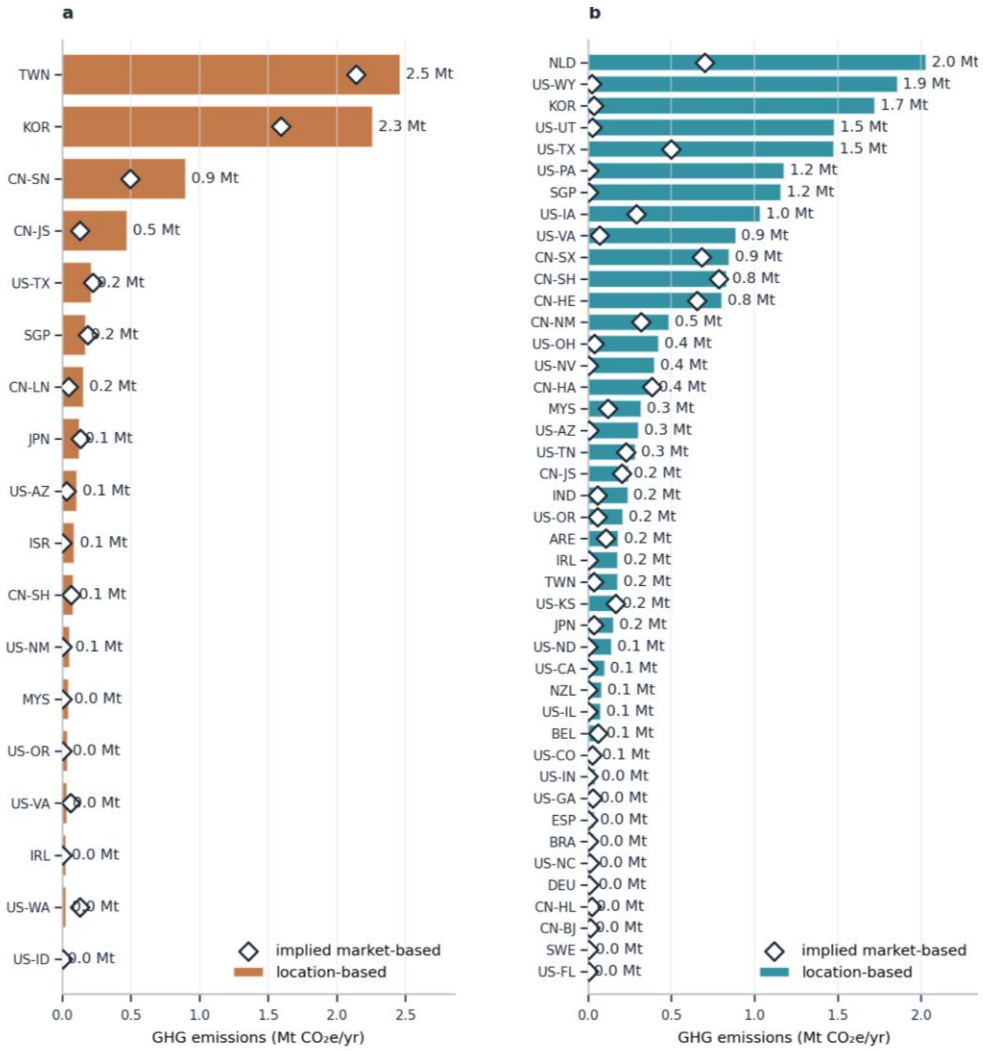

507

508 ***Extended Data Fig. 2 | Exact-geography emissions rankings.*** *a, Manufacturing emissions*
509 *ranked by exact geography under the location-based basis, with diamonds showing market-*
510 *based values applied to the same activity weights. b, Operations emissions for resolved data-*
511 *center geographies under the same comparison. Market-based values use the approved facility-*
512 *or company-boundary factors from the evidence protocol and should not be read as facility-*
513 *specific renewable electricity claims unless supported by the relevant evidence tier.*

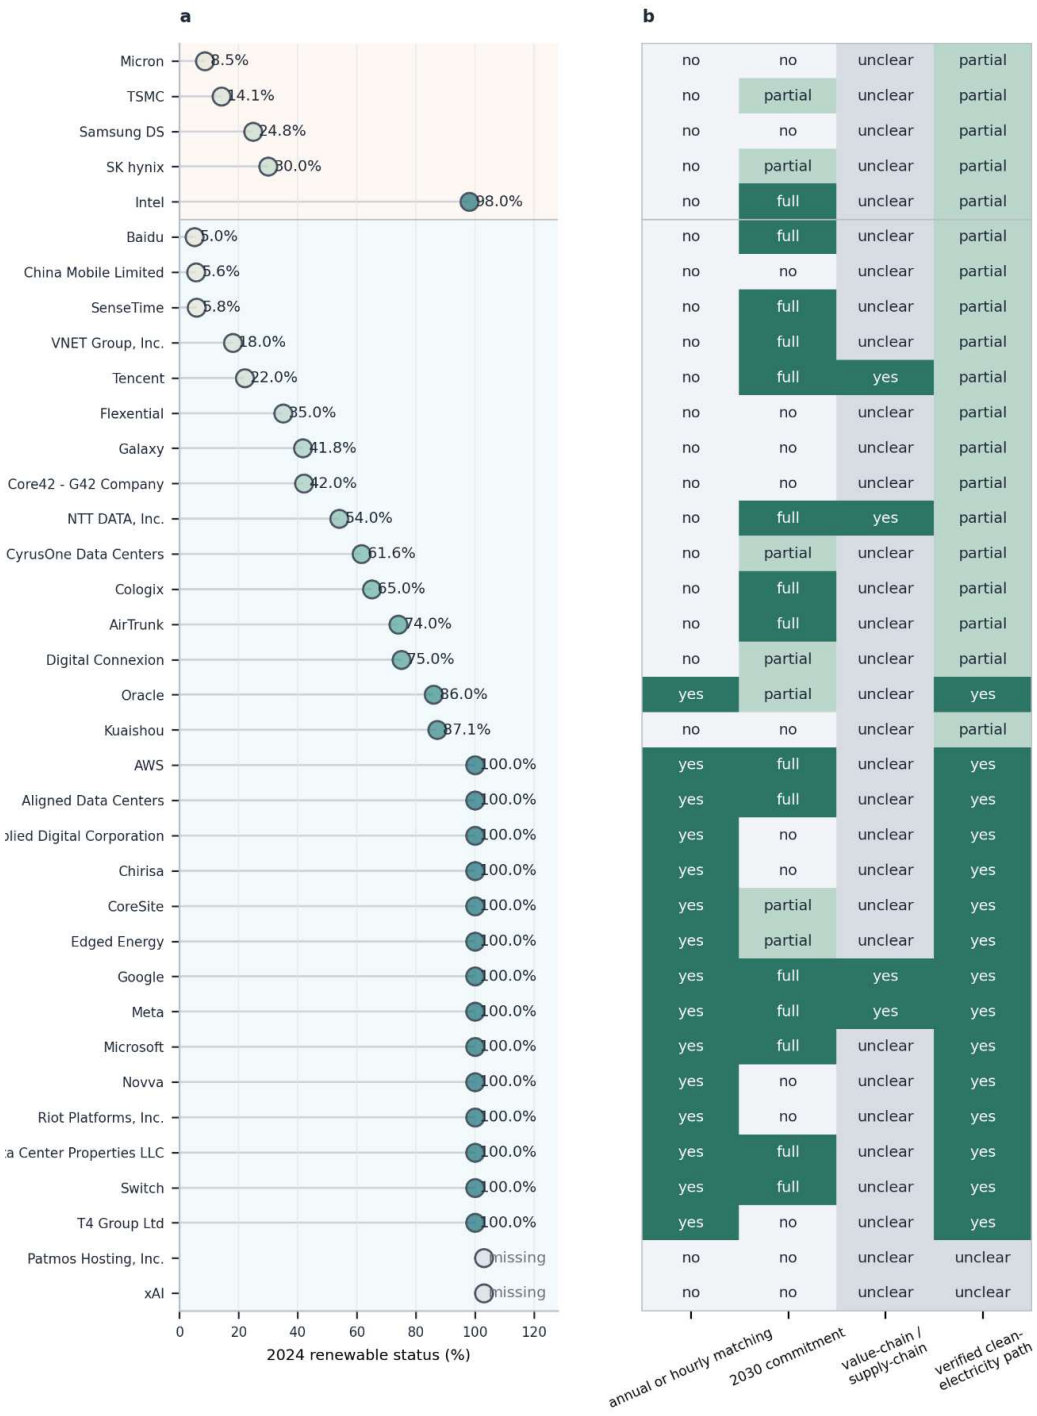

514

515 ***Extended Data Fig. 3 | Renewable electricity status and maturity.*** *a, Reported 2024 renewable*

516 *electricity status for manufacturing suppliers and operations entities; missing values are*

517 *retained as missing. b, Qualitative maturity indicators for annual or hourly matching, 2030*

518 *commitments, value-chain or supply-chain action and verified clean-electricity pathways. These*

indicators support source-policy interpretation and do not change central location-based or market-based estimates unless they provide a comparable, attributable market-based factor.

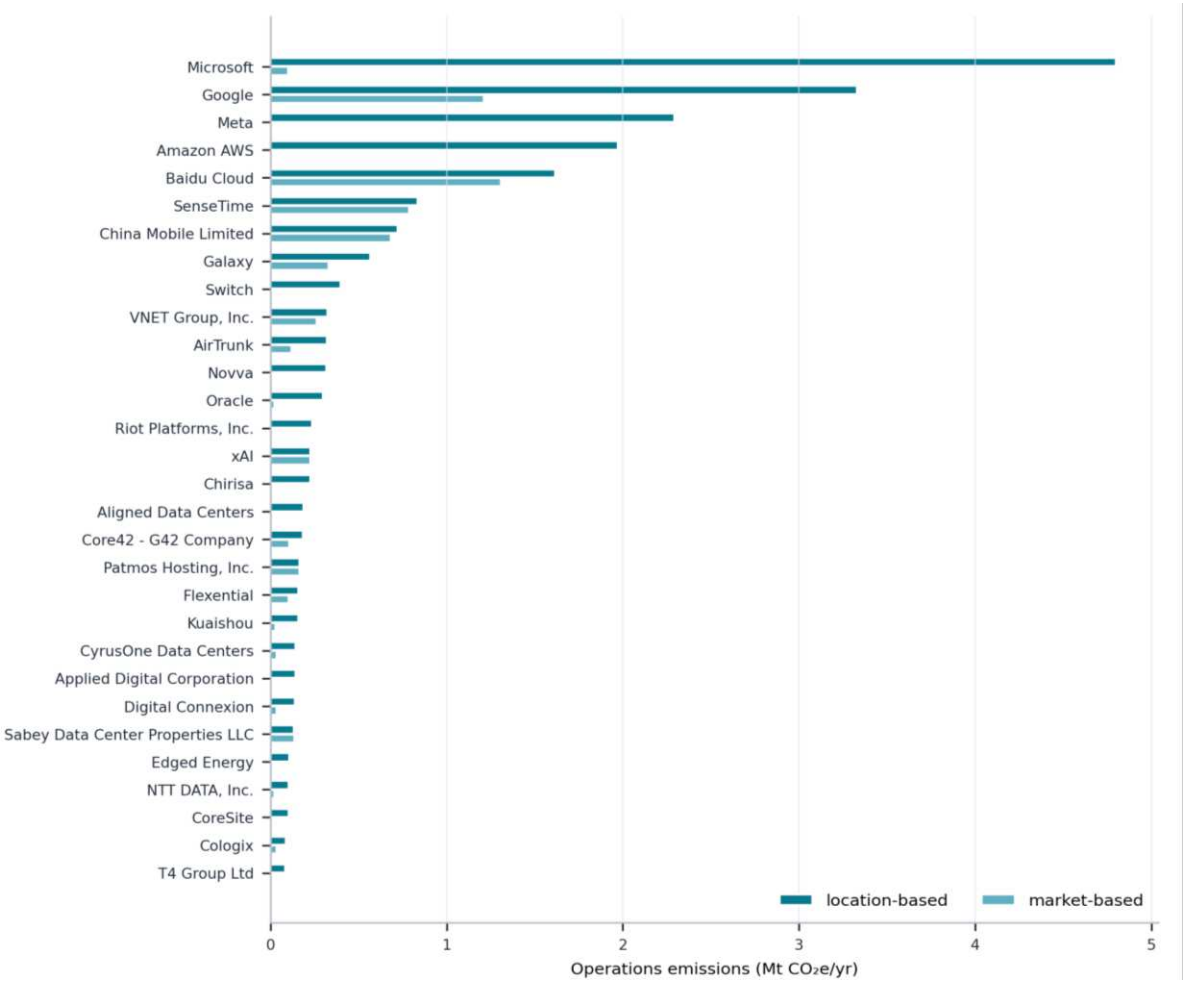

Extended Data Fig. 4 | Operations emissions by company. *Bars aggregate approved facility-source records by matched company. Location-based emissions use local or regional emission factors; market-based emissions use approved factors selected through the tiered evidence protocol. Facility electricity is held fixed, and records without a clear attributable market-based emission factor retain the location-based factor.*

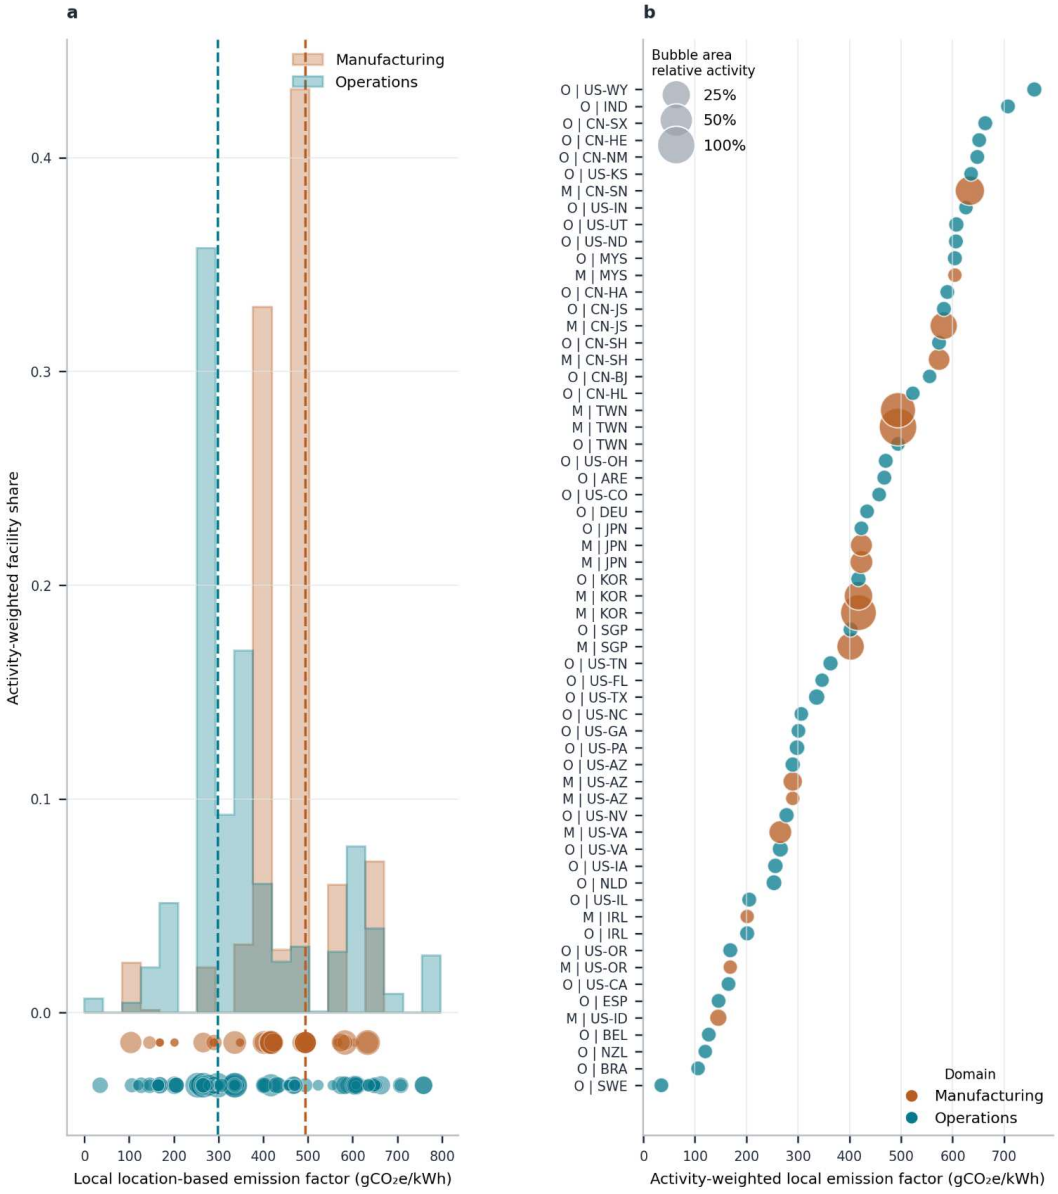

527

528 ***Extended Data Fig. 5 | Emission-factor exposure distributions.*** *a, Activity-weighted local*

529 *location-based emission-factor distributions for manufacturing and operations, with dashed lines*

530 *showing activity-weighted central tendencies. b, Sorted exact-geography exposure audit for*

531 *manufacturing and operations. The figure tests whether location-based results are driven by a*

532 *few high-emission-factor geographies or by broader activity-weighted exposure.*

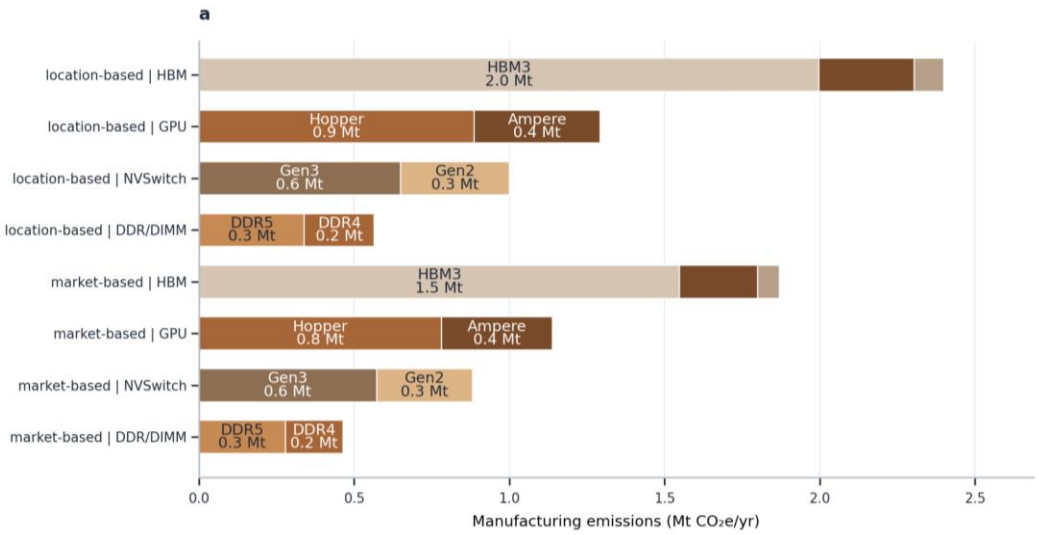

Extended Data Fig. 6 | Hardware product-family decomposition. *Manufacturing emissions by product-family layer for selected hardware module groups under location-based and market-based accounting. The panel supports module-level interpretation while preserving the 12-module boundary and distinguishing assigned fractions from the unallocated remainder.*

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformation.docx
- Supplementarydataappendix.xlsx
- ExtendedDataFig.2.pdf
- ExtendedDataFig.5.pdf
- ExtendedDataFig.1.pdf
- ExtendedDataFig.3.pdf
- SIGuide.doc
- ExtendedDataFig.6.pdf
- ExtendedDataFig.4.pdf