

How Microsoft is advancing embodied carbon measurement at scale for Azure hardware

By: Kali Frost, Ines Sousa and Leo Aspauza

Contributors: Arash Saboori, Lalit Joshi, Ranganathan Srikanth, Daniel Berger, Fiodar Kazhamiaka



Executive Summary

This white paper introduces Microsoft's cloud hardware emissions methodology (CHEM), a scalable approach to measure the embodied carbon of datacenter information and communication (ICT) hardware. CHEM combines process-based life cycle assessment (pLCA), AI-powered cloud tools, product granular data and state-of-the-art semiconductor impact data that enable fleet-wide carbon accounting and actionable insights for supply chain interventions, hardware design, and sustainability reporting.

What this white paper covers:

- The challenge of accurately measuring embodied carbon at scale in datacenter hardware and why traditional LCA approaches are limited.
- How CHEM works: integrating Microsoft's product data systems, automated material and proxy mapping, and advanced semiconductor models.
- Applications across Microsoft's business: Scope 3 reporting, supply chain interventions, hardware design optimization, and carbon roadmapping.

CHEM shifted Microsoft from financial proxies to pLCA, empowering sourcing and design teams with actionable measurements and ability to cost-effectively prioritize interventions. AI amplified this impact by automating steps, such as materials to life cycle inventory (LCI) mapping, reducing modeling time, and enabling LCA practitioners to focus on integrating supplier data and driving actionable insights.

Key benefits of CHEM that are highlighted include:

- Scalability and automation: CHEM models thousands of hardware configurations efficiently, enabling rapid, consistent, and high-resolution carbon footprinting across Microsoft's cloud fleet.
- Data quality and modularity: Integration of supplier-specific data and state-of-the-art semiconductor impact data improves accuracy, while modular architecture allows for continuous updates as new data becomes available.
- Actionable insights: CHEM enables hot spotting emissions drivers deep within multi-tiered supply chains, supporting targeted decarbonization interventions and informed hardware design decisions.

We also emphasize key areas for next-level data-driven actions in the industry, which include invest in better data, use of AI to improve data exchange and quality, quantify boundaries differences between LCA methodologies, and collaborate to standardize LCA based carbon accounting for data center ICT hardware.

CHEM's impact extends beyond Microsoft. As part of ongoing collaboration with other hyperscalers and the broader ICT sector to share best practices and align on unified standards, CHEM is being leveraged, alongside partner methodologies, to develop Product Category Rules (PCRs) and an open methodology to scale LCA. These efforts, driven through the Open Compute Project (OCP) and SEMI Semiconductor Climate Consortium (SCC), aim to harmonize carbon accounting practices, improve data quality, and enable consistent, actionable measurement of embodied carbon emissions that supports global sustainability goals.

Contents

Executive Summary	2
Introduction	4
Why CHEM?	5
CHEM: Advancing LCA for Cloud	7
CHEM in Action	10
Conclusion and What's Next	13
References	14
Glossary	15

Introduction

Microsoft and other large-scale cloud service providers are targeting significant greenhouse gas (GHG) emissions reductions to align with net zero global goals. To meet these goals, it is critical to measure and reduce the Scope 3 embodied carbon impact of Information and Communication Technology (ICT) hardware across datacenters and requires a reliable, granular, and scalable process for calculating GHG emissions and tracking progress. The datacenter ICT hardware supply chain is one of the largest contributors to Microsoft's Scope 3 GHG emissions, thus actionable measurement of ICT hardware embodied carbon is required to accelerate decarbonization efforts and meet our 2030 carbon negative goal.

This white paper introduces Microsoft's cloud hardware emissions methodology (CHEM), a life-cycle assessment (LCA) based methodology developed to automate and scale the environmental impact modeling of complex ICT hardware systems in our datacenters. CHEM addresses key challenges that have hindered the use of LCA for these systems:

Scaling product LCA to a datacenter fleet

LCA provides the granularity required to identify hotspots driving emissions, quantify impact reduction from implementing targeted interventions, and track decarbonization progress towards goals – CHEM modeling helps us scale these insights across our Azure fleet without compromising on precision.

Rapid technological development

Automation frees up time to focus where it matters – CHEM modeling automation enables effort to be applied in improving data quality and supply chain representativeness, in an industry where systems are rapidly evolving and materials are only increasing in complexity.

Integration of semiconductor impact data

Semiconductor chips are key drivers of datacenter ICT embodied carbon – CHEM's ability to differentiate impact among key chip technologies (e.g. memory, processors) with state-of-the-art impact data in product and fleet models can improve decision-making for supply chain decarbonization and cloud systems architecture design.

We provide insights into CHEM and how it has been applied to Microsoft's cloud business. We highlight a range of key applications, including organizational Scope 3 emissions accounting and reporting, supply chain decarbonization interventions, ICT hardware systems design for sustainability, and carbon roadmapping. We also emphasize in concluding remarks key areas we believe should be in the spotlight for next-level data-driven actions in the industry to accelerate the decarbonization of cloud infrastructure and contribute to global sustainability goals.

Why CHEM?

Microsoft and other major cloud providers are targeting significant GHG emissions reductions to align with net zero global sustainability goals [1]. For Microsoft to drive its decarbonization efforts and meet its carbon negative goals by 2030 [2], it is critical to measure the Scope 3 embodied carbon impact of ICT hardware across its data centers accurately, actionably, and at scale. Additionally, the embodied carbon in semiconductor chips represents the main contributor to Microsoft's datacenter ICT hardware supply chain emissions. Therefore, integrating state-of-the-art semiconductor environmental impact data is required to track this rapidly evolving industry.

Assessing Scope 3 embodied emissions for datacenter ICT hardware presents significant challenges but remains an achievable objective. The complexity and large-scale nature of these technologies—combined with multi-tiered supply chains, rapid technological advancement, and a globally distributed supplier network—introduce substantial obstacles to accurate measurement and tracking of embodied carbon emissions. These challenges are further compounded by limited visibility into supplier operations, competitive pressures that restrict transparency, and inconsistencies in carbon accounting methodologies. Multiple strategies must be employed to tackle this massive challenge, including supplier engagement, systematic data collection and enhancement of data quality, as well as harmonization of methodologies and industry-wide collaboration [3]. Further, the ability to scale carbon accounting of complex datacenter ICT hardware systems beyond a few products, while maintaining the data and model granularity required for implementing and tracking decarbonization interventions, is necessary to enable actionable carbon impact measurement and tracking.

Life-cycle assessment (LCA) is an established and internationally standardized framework to systematically quantify the life-cycle environmental impacts of products, processes or services, including their “cradle-to-gate” embodied carbon [4]. In a top-down approach, economic input-output (EIO) based LCA that links economic activity to environmental impacts can be applied to directionally and comprehensively estimate embodied carbon impacts [5]. However, common limitations of EIO LCA, particularly lack of disaggregated product-specific data and economic-based inventory estimates, make it difficult to account for rapidly evolving and innovative ICT hardware assets or model specific technical process-based interventions to quantify impacts of IT hardware system design changes. In contrast, process-based LCA (pLCA) can help identify product, material, or process-specific hot spots and measure and track targeted modeled emissions reductions – particularly when high-quality unit-process engineering inventory and variances are available. This level of data granularity and parameterization can support programs driving business decarbonization efforts, such as supplier engagement and eco-design.

Although pLCA is widely regarded as a powerful methodology, it faces limitations [6] that can hinder its broader application—especially given the fast-paced evolution of technology and the complex supply chains involved in datacenter ICT hardware manufacturing. Below, we outline some of the key challenges.

Reliance on generic models and secondary data

Generic life-cycle inventory (LCI) process data provided by public or commercially available databases are heavily used to model impacts. While LCI datasets can provide a solid baseline, processes are typically modeled based on industry averages and/or literature data. For the ICT sector, in particular, these processes are often outdated and do not reflect the latest technologies or company-specific supply chains.

Some commercial models provide impacts that are aggregated across many materials, manufacturing processes, or tiers of the supply chain, preventing the LCA practitioner from incorporating company-specific supply chain data, for either baseline estimation or at the level of a decarbonization intervention (e.g. carbon free energy use at a chip manufacturing facility).

Limited scalability

Data collection is often conducted manually through product teardowns to create a representative Bill-of-Materials (BOM) for a part, subassembly or assembly. And even if a BOM is already well characterized, it must often manually be entered into the LCA modeling software. Additionally, one of the most time-consuming steps is where the LCA expert needs to associate each material in the BOM with the most representative materials and manufacturing processes in the LCI database. These processes can take up to 100+ hours per server (>90% of effort based on an estimate of the time taken to map materials to LCI processes), which limits the ability to scale the use of LCA across hundreds or thousands of unique products, without compromising on representativeness.

Datacenters can house several thousands of server and rack configurations (Figure 1). To measure the embodied carbon within our ICT hardware infrastructure and be able to design and track interventions over time, we need automated, scalable methodologies that deliver consistent results across materials, components, subassemblies, assemblies, and server or rack configurations. This comprehensive approach is essential for ensuring thorough coverage and minimizing modeling time.




Figure 1. Generic example to illustrate the complexity and scale of ICT hardware in datacenters

CHEM: Advancing LCA for Cloud

CHEM utilizes an innovative approach for conducting LCA to automate and scale the environmental impact modeling of complex datacenter hardware systems. It was developed using a combination of Azure data services, a cloud-based automated LCA software (Makersite) and a proxy mapping tool that incorporate Artificial Intelligence (AI) and state-of-the-art semiconductor manufacturing LCIs developed by the imec Sustainable Semiconductor Technologies and Systems (SSTS) program¹.

CHEM at a Glance

CHEM incorporates three key innovation areas which may differ from common practice (see Figure 2):

- 1. Integration of Microsoft's product data management (PDM) system and full materials declaration (FMD) data with the LCA software:** Through API integration with internally managed systems, we have streamlined the overall process of transferring BOM hierarchies and associated material data into the LCA modeling environment. These robust internal data management systems are managed by dedicated teams that collaborate closely with hardware design teams and suppliers to collect and maintain data for thousands of components. This integration has not only reduced the effort required to enrich LCA models but also decreased the risk of errors and enhanced modelers' ability to update assessments in sync with hardware changes.
- 2. Automated assignment of a material to the appropriate life-cycle inventory (LCI):** The LCA software automates the process of associating the specific material composition of a part (i.e. FMD) with highly representative unit processes available from third-party reviewed LCI databases, including ecoinvent [7]. This cuts out a lot of manual effort and, while LCA experts are still involved in the process, they can now focus on integrating suppliers' primary data, performing data quality analysis, ensuring the model is representative, and providing actionable insights for sourcing teams and designers. For the semiconductor materials in our high-impact chips, we follow a specific, semi-automated process to help collect custom information and integrate data from the imec SSTS program.

¹ imec (Interuniversity Microelectronics Center) is an R&D center headquartered in Belgium, which specializes in creating technology roadmaps for the semiconductor industry. The Sustainable Semiconductor Technology and Systems (SSTS) program was launched in 2021 to update and improve environmental impact assessment of semiconductor manufacturing.

3. **Scaling granular LCA data across Azure hardware:** To scale to thousands of configurations representing the Azure hardware fleet, it requires post-processing using Azure data services. Key steps in this process are:

- Creation of a CHEM parts database. To speed up processing of our high-complexity BOMs, LCAs are completed for each part, which then appear repeatedly across rack BOMs within our datacenters.
- Tools were built to semi-automate the selection of part proxies, by providing a proxy part recommendation to the LCA practitioner. For parts where complete materials data is not available, a quality proxy must be selected from the parts database. Proxy selection is based on a similarity algorithm to identify the likeness between parts based on key part characteristics. For example, an SSD may have a few fundamental specifications which account for most of its emissions – such as the semiconductor technology node, NAND chip storage capacity, and manufacturing location. LCA practitioners can then quickly select from available quality proxies.
- Automated reconstruction of the PDM BOM hierarchy with CHEM emissions appended. We automate the process of quantifying the emissions of each unique server or rack configuration by using the CHEM parts database and re-constructing the server or rack emissions by leveraging the BOM structure within our existing PDM system.

LCA practitioners perform data quality checks to ensure that inputs to the LCA software (BOM and FMD) and subsequent data transformations are accurate and representative.




Figure 2. Generic examples to illustrate CHEM.

Customizing Chip LCA Models with imec SSTS Data

For Microsoft's cloud ICT hardware, our LCA models show that semiconductor chips are the single most important driver of embodied carbon. To reduce the impact of this key hotspot, high quality and transparent data is needed to support prioritization and decision-making, as well as to ensure accurate emissions reporting. CHEM incorporates imec.netzero data [8] from the imec SSTS program to assess the environmental impact of manufacturing current and future logic, DRAM and NAND in high-impact parts (e.g., CPUs, GPUs, DIMMs and SSDs). These parts constitute ~70-85% of a server's impact, depending on the hardware configuration.

The virtual fab models in imec.netzero include detailed process flows enumerated by imec experts with data collected from imec's 300 mm fab and associated ecosystem of material and equipment providers. This state-of-the-art LCI data for semiconductor manufacturing is used to create custom unit process models that are technology and node specific, for example, a logic processor manufactured using the 7nm process node. We create custom models in Makersite's product lifecycle intelligence platform by incorporating the imec.netzero LCI data into existing ecoinvent based LCA models.

We further tailor our LCA models by incorporating product- and supply chain-specific data. This includes aligning the electricity grid mix with the geographic location of the semiconductor manufacturing facility and specifying chip-level parameters, such as technology node, die size, and yield estimates. These enhancements allow us to generate LCAs that more accurately reflect actual ICT hardware configurations in our datacenters. Additionally, the integration of supplier-provided FMDs for IC (chip) packaging ensures that the material composition is precisely represented, thereby improving the fidelity of our data.

CHEM Benefits

CHEM provides key benefits, including:

Enables scalability: It allows parts, assemblies, and configurations across Microsoft's cloud hardware fleet to be modeled and managed.

Increases modularity that supports model customization to Microsoft's cloud ICT hardware supply chain: The modular architecture of fully disaggregated LCA models, mapped to the BOM and hardware's chemical composition, allows us to replace generic component and manufacturing process data with supplier-specific data as it becomes available, which in turn propagates across the models.

Improves consistency between LCA practitioners in modeling and impact dataset selection:

Automated mapping helps standardize modeling decisions, minimizing inconsistencies and human error that may arise from individual LCA practitioners, for example, in selecting ecoinvent datasets.

Speeds up the LCA modeling process: With less manual intervention, LCA practitioners, rather than spending hours manually mapping datasets, can allocate that time to gathering and processing primary data from suppliers and performing data quality assurance and analysis that support decision making by sourcing and eco-design teams.

The use of AI in this approach to automate processes is an essential time-saving component, and makes its implementation, at scale, possible. It enables LCA practitioners to spend more time focused on communicating insights, identifying key decarbonization efforts, and integrating results into

comprehensive roadmapping efforts. However, we want to stress the importance of a human-in-the-loop for quality assuring model results, and for prioritizing and improving model performance.

CHEM in Action

CHEM has been used in different functional areas of Microsoft's business. The scalability and process-level modeling underlying CHEM increased ICT hardware coverage and model resolution across a range of key applications to Microsoft's cloud business, including:

- Organizational Scope 3 emissions accounting and reporting,
- Supply chain decarbonization interventions,
- ICT hardware systems design for sustainability, and
- Actionable carbon roadmapping.

Scope 3 Emissions Reporting

To account for the embodied carbon within our datacenter fleet—information that is incorporated into Microsoft's annual sustainability report and other external disclosures—the application of CHEM has enabled coverage of over 97% of cloud server racks using highly granular pLCA based data. Additionally, nearly 80% of semiconductor emissions are now quantified using imec.netzero data, a level of detail and scale that would otherwise require thousands of hours of manual analysis to achieve.

Supply Chain Interventions

The supply chain of Azure hardware remains one of Microsoft's key levers for decarbonization. A prominent use case of product LCA is to identify key intervention hotspots which can often occur from 1-5 tiers deep within the supply chain, based on ownership structure and specialization of the supply chain. While this is nothing new for product LCA, our ability to aggregate this information across suppliers, parts, assemblies or programs is enhanced by the coverage and scalability of CHEM. For example (Figure 3), if we can pinpoint the impact of electricity or direct emissions in a given process, there are opportunities to scale this quantification, size and advocate for the targeted use of carbon free energy or manufacturing process interventions that reduce direct emissions, and account for this in our inventory with high precision.




Figure 3. Targeting and sizing supply chain interventions with CHEM. In this generic example, we can trace the impact of an electricity hotspot in the Tier 3 supply chain (“Raw Materials Processing”) by (1) material linkage: identify the quantity of material sourced from the component’s FMD, (2) LCI association: connect to material’s LCI dataset, and (3) impact scaling: calculate how the impact cascades to the fleet level.

Hardware Systems Architecture Design

Major advances in procurement of carbon free energy have resulted in more awareness within the ICT systems architecture community and shifting from primarily focusing on operational energy efficiency, to also better understanding the drivers of embodied emissions [9]. CHEM LCA data has been used by system architects at Microsoft to quantify the impact of server design on datacenter’s carbon emissions, providing a reliable way to integrate carbon metrics in the design of a server, rack, or cluster [10]. This has also led to an increased confidence in the magnitude of emissions and the hotspots that are key from a cloud systems-level (operational + embodied) design perspective and contributed to a deeper understanding of the impact of datacenter design optimizations [11]. Aggregating emissions at the datacenter fleet level has led to key insights regarding the impact of storage and memory (Figure 4), also noted by others in the industry [12].




Figure 4. Key contributors to the embodied Scope 3 emissions impact of the Azure fleet hardware. Memory and storage chips include DRAM and NAND. Processor chips include those present in CPUs, GPUs, ASICs, and FPGAs.

Carbon Roadmapping

Supply chain interventions, ICT hardware design optimizations, software efficiency initiatives and other carbon reduction interventions are compiled to create our Scope 3 carbon reduction roadmap – i.e. the decarbonization glide path needed for cloud ICT hardware's contribution towards Microsoft's scope 3 emissions targets. CHEM is now the methodology underlying the development and update of this roadmap which reflects and reports interventions at component, rack, and programmatic levels, offering increased resolution and accuracy in targeting interventions and tracking progress towards our goals.

Importantly, CHEM pivoted us from simplistic financial-based approaches to process-based modeling of our cloud hardware fleet, which impacted our roadmapping and interventions in two key ways:

- Enhanced with imec.netzero data, CHEM led to a deepened understanding of the impact of memory, which informed our supplier engagement and eco-design focus and efforts in decarbonization.
- This 'telemetry' gave us a more precise estimate of where embodied carbon impacts reside, allowing more effective return on investment modeling associated with our internal carbon fee.

Conclusion and What's Next

Enhancing the precision and scale of actionable embodied carbon data for cloud ICT hardware is essential to achieving our 2030 carbon reduction goals. By improving the granularity and scalability of carbon footprint assessments—incorporating process-based LCA data throughout every level of the BOM for nearly all ICT hardware in datacenters—we can more effectively target the hotspots, size and prioritize decarbonization interventions, drive supply chain action, and support carbon-conscious product design decisions.

CHEM brings process-level impact modeling resolution, at scale, with an innovative approach and tooling for conducting LCA that automates and scales carbon impact modeling of complex data center ICT hardware systems. Developed using a combination of Azure data services, cloud-based LCA software and a proxy mapping tool that incorporate AI, and imec's state-of-the-art semiconductor environmental impact data, CHEM has enabled Microsoft to take a leap in actionable measurement of carbon emissions for our company. Specific business applications were described in this white paper.

To build on these efforts, there are several key areas we believe should be in the spotlight for next-level data-driven actions in the industry to accelerate the decarbonization of cloud infrastructure and contribute to global sustainability goals:

- Continue investing in high-quality LCI data generation, as the robustness of an LCA depends on the underlying data quality. Prioritize improvements based on potential impact – focusing on enhancing third-party datasets for Tier2+ processes (e.g., semiconductor-grade chemicals production) and incorporating more high-quality primary data from suppliers for identified hotspots.
- Explore the potential of AI and machine learning to enhance data exchange, data quality, and further optimize system design for sustainability.
- Understand and quantify the relative impact and accordingly harmonize and account for boundary differences between pLCA, EIO LCA, and organizational carbon footprinting.
- Standardize carbon footprinting for datacenter ICT hardware. Microsoft is collaborating with industry partners through the Open Compute Project (OCP)² and SEMI SCC³ to develop Product Category Rules (PCRs) for datacenter equipment and semiconductors. These efforts aim to harmonize LCA-based carbon accounting, enabling consistent and actionable measurement of embodied emissions and impact. As part of these ongoing standardization efforts, we are also working with other hyperscalers to develop and share with the broader industry a unified methodology [13], leveraging each company's best practices for scaling LCA.

² OCP (Open Compute Project) is a nonprofit foundation that promotes the open sharing of designs for data center hardware and infrastructure. Sustainability was introduced as OCP's fifth tenet in 2022, making it a core principle alongside openness, efficiency, and scalability.

³ SEMI SCC (Semiconductor Climate Consortium) is a collaborative industry platform under SEMI (Semiconductor Equipment and Materials International) for companies across the semiconductor value chain to collectively set and address decarbonization of the sector.

References

[1] Science Based Targets initiative (SBTi), "Companies Taking Action," Target Dashboard [Online]. Available: <https://sciencebasedtargets.org/target-dashboard>

[2] Microsoft, "2025 Environmental Sustainability Report," May 29, 2025. [Online]. Available: <https://www.microsoft.com/en-us/corporate-responsibility/sustainability/report/?msocid=02e4e384d6b668bc29d5f7e8d76f6908>

[3] McKinsey & Company, "Making Supply-Chain Decarbonization Happen," June 14, 2021. [Online]. Available: <https://www.mckinsey.com/capabilities/operations/our-insights/making-supply-chain-decarbonization-happen>

[4] International Organization for Standardization (ISO), "ISO 14040:2006 — Environmental management — Life cycle assessment — Principles and framework." [Online]. Available: <https://www.iso.org/standard/37456.html>; "ISO 14044:2006 — Environmental management — Life cycle assessment — Requirements and guidelines". [Online]. Available: <https://www.iso.org/standard/38498.html>

[5] C. T. Hendrickson, L. B. Lave, and H. S. Matthews, *Environmental Life Cycle Assessment of Goods and Services: An Input-Output Approach*, 1st ed. New York, NY, USA: Routledge, 2006.

[6] Microsoft, "Microsoft Consumer Devices Life Cycle Assessment Methodology—Executive Summary (v2.1)," April 2024. [Online]. Available: <https://www.microsoft.com/en-us/download/details.aspx?id=55974&mocid=274140e96e1d650a34c356646f1e64c2>

[7] R. Frischknecht, N. Jungbluth, H.-J. Althaus, *et al.*, "The ecoinvent Database: Overview and Methodological Framework," *Int. J. Life Cycle Assessment*, vol. 10, pp. 3–9, 2005.

[8] imec, imec.netzero virtual fab. [Online]. Available: <https://netzero.imec-int.com>

[9] D. Berger, D. Brooks, F. Kazhamiaka, *et al.*, "Reducing Embodied Carbon is Important," ACM SIGARCH Blog, Aug. 3, 2023. [Online]. Available: <https://www.sigarch.org/reducing-embodied-carbon-is-important/>

[10] J. Wang, D. S. Berger, F. Kazhamiaka, *et al.*, "Designing Cloud Servers for Lower Carbon," ISCA 2024. [Online]. Available: <https://www.microsoft.com/en-us/research/wp-content/uploads/2024/03/2024-GreenSKU-ISCA2024.pdf>

[11] H. Alissa, T. Nick, A. Raniwala, *et al.*, "Using Life Cycle Assessment to Drive Innovation for Sustainable Cool Clouds," *Nature*, vol. 641, pp. 331–338, Apr. 30, 2025. [Online]. Available: <https://link.springer.com/article/10.1038/s41586-025-08832-3>

[12] L. Rivalin, L. Yi, M. Diefenbach, *et al.*, "Estimating Embodied Carbon in Data Center Hardware, Down to the Individual Screws," Sep. 10, 2024. [Online]. Available: <https://sustainability.atmeta.com/blog/2024/09/10/estimating-embodied-carbon-in-data-center-hardware-down-to-the-individual-screws/>

[13] L. Rivalin, S. Weiss-lijn, A. Saboori, "From Component to Fleet: An Open-Sourced Methodology for IT Carbon Accounting at Scale," OCP Global Summit, Oct. 2025. [Online]. Available: https://drive.google.com/file/d/1Ohx4vgneKeVqeTl-zY_Fg7_tMXjvA4pd/view

Glossary

300 mm Fab: fabs produce chips used in servers and devices and are often described by the size of wafer they produce. A 300 mm or 12-inch wafer fab produces a 300 mm diameter IC wafer, which is then cut into individual die (chips).

ASIC (Application-Specific Integrated Circuit): A custom-designed chip for a particular application.

Bill of Materials (BOM): A hierarchical structure of parts, components, and assemblies that make up a hardware asset such as a rack assembly.

Central processing unit (CPU): One high impact type of component found in server, commonly considered the "workhorse" of a server.

Cradle-to-Gate: Assessment of a product's environmental impact from raw material extraction (cradle) to the point it leaves the manufacturer (gate).

Eco-design: Incorporation of environmental aspects into the product development process, to minimize environmental impact throughout the product life cycle.

Die area: A quantity property representing the total area of an integrated circuit die.

Dual in-line Memory Module (DIMM): Module that contains one or several random access memory (RAM) chips on a small circuit board with pins that connect it to the motherboard.

Fleet: In cloud hardware context, it is the collective set of hardware assets (servers, racks, components) deployed across datacenters

FPGA (Field-Programmable Gate Array): A reconfigurable integrated circuit used for specialized computing tasks.

Full material disclosure (FMD): A file representing the full material breakdown by mass of a part, originally intended for compliance purposes.

Graphics Processing Unit (GPU): One high impact type of component found in server, designed to handle tasks that require high speeds and parallel processing.

Greenhouse Gas (GHG): Gases in the atmosphere absorb radiation, trap heat, causing the greenhouse effect making the planet warmer. The seven GHGs listed in the Kyoto protocol and the GHG protocol corporate accounting and reporting standard are: carbon dioxide (CO₂), perfluorocarbons (PFCs), sulfur hexafluoride (SF₆), methane (CH₄), nitrous oxide (N₂O), and hydrofluorocarbons (HFCs), and nitrogen trifluoride (NF₃). The most significant emissions come from human activities from burning fossil fuels for electricity, heat, and transportation.

High Impact Parts (HIPs): Parts that are known to have higher environmental impacts compared to other components.

Hyperscaler: Large-scale cloud service provider that operates globally distributed datacenters and delivers computing, storage, and networking resources at scale.

Integrated Circuit (IC): A set of electronic circuits on a single silicon chip.

Life Cycle Assessment (LCA): A standardized approach to quantify the environmental impacts of a product or service across its life cycle, from raw materials extraction, production, use, end-of life and includes transportation through all stages.

Life Cycle Inventory (LCI): Detailed inventory of manufacturing processes collected during an LCA study that lists all the inputs of materials and energy and the outputs of products, waste and emissions.

Memory Module (MMOD): A high impact component found in a server that is comprised of a printed circuit board with large quantities of DRAM (dynamic random-access memory).

NAND Flash Memory: A type of non-volatile storage technology that does not require power to retain data.

Primary Data: Data collected directly from specific processes, suppliers, or operations, typically considered more accurate and representative.

Product Category Rules (PCR): Set of standardized rules that define how to conduct an LCA for a specific category of products to ensure consistency and comparability (e.g. system boundaries, functional unit, data quality and sources, allocation rules). They are aligned with ISO standards (14040/44/67) and developed through multi-stakeholder consultation to ensure transparency and industry alignment.

Product Teardown: The disassembly of a product and its parts to describe its construction, attributes, and functionalities.

Rack Assembly: A large cabinet that houses servers or other networking, compute, or storage, and power handling equipment used in Microsoft datacenters.

Scope 1, Scope 2, Scope 3 GHG Emissions: Scope 1: Direct GHG emissions from owned or controlled sources (e.g., company vehicles, on-site fuel combustion); Scope 2: Indirect GHG emissions from the generation of purchased electricity, steam, heating and cooling. Scope 3: All other indirect emissions that occur in a company's value chain, including upstream and downstream activities (e.g., supplier manufacturing, product use, transportation).

Solid State Drive (SSD): A data storage device that reads and writes to non-volatile solid-state memory instead of rotating magnetic platters for data storage. SSDs have been identified as one of the high impact components found in a server.

Stock keeping unit (SKU): A retailer's product number that tracks stock of merchandise. In Microsoft cloud hardware and this document, a SKU represents a hardware program and often represents the entire BOM, from the rack assembly (RASSY) level down to the components.

Technology Node: Also referred to as tech node, process node, or simply node, a specific semiconductor manufacturing process, often described in units of length (nanometers), that refers to half the distance between identical features. Nodes are known to get smaller over time.

Technology Type: General terms used to segment integrated circuit products. The three main technology types referred to in this document are Logic, DRAM (dynamic random access memory), and NAND (i.e., flash memory).

Unit Process: In life cycle assessment (ISO 14040): the smallest element considered in the life cycle inventory analysis for which input and output data are quantified.

This white paper is for informational purposes only. Microsoft makes no warranties, express or implied, in this document.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2026 Microsoft Corporation. All rights reserved.