



Solving the *AI paradox*: Sustainability and business value in enterprise deployments

How organizations can unlock
value and prioritize efficiency

The challenge of balancing the dual impacts of AI

What's the true environmental impact of the latest revolution in AI? Many organizations, from those building models to those deploying them, are struggling to quantify it¹. Generative AI and large language models have embedded themselves in business processes, and even daily life, faster than anyone anticipated. And with every new iteration, its potential grows.

But this growth has a significant downside. As hyperscalers continue to expand their data centers and model catalogs, and more organizations strive

to deploy production-ready AI, the GHG emissions, resource use, land requirements, compute capacity demands, and costs are becoming unavoidable.

With the right planning and mitigation techniques, it's possible for leaders to take a more sustainable approach. This will allow organizations to pursue their AI ambitions while managing their footprint.

This guide unpacks the paradox of AI's potential to advance sustainability goals versus its impact on the environment, and explores ways Capgemini can help leaders make the right decisions to maximize its value.

The complexity of the AI value chain creates gaps in leaders' knowledge

For organizations in every sector, balancing the costs of AI with the outcomes they can achieve is a vital but often overlooked step in planning and executing initiatives. In the industry's rush to embrace AI, efficient, optimized usage is coming second to exciting capabilities and big-ticket productivity gains. It's time to change that.

Though it's narrowing, there's still an environmental awareness gap in many leadership teams, and it's one of the key efficiency roadblocks. Just 12% of organizations currently measure the impact of their generative models². That's often because there are so many layers in the AI production and value chain. Through hardware manufacturing, model development, and into deployment, a lot of stakeholders are sharing responsibility for emissions and resource use.

If there are any unclear processes or missing data points at any stage, this lack of clarity rapidly compounds, making it far more difficult for organizations at the end of the chain to measure and manage the total footprint of their generative AI projects.

This lack of visibility may be due to:

Difficulty measuring effectively: The simplest answer is that emissions and usage can be hard to monitor, especially without industry-wide standards for exactly what should be measured and how to report on it.

A lack of reporting: If just one link in the chain, whether it's a mining facility, component manufacturer, or data center operator, fails to measure and share its footprint, a crucial part of the overall picture will be missing.

Limited transparency: AI model providers are often opaque about the footprint of their models, potentially due to the risk of exposing sensitive information about their proprietary technologies. Unlike open-source models, they can't be independently evaluated by researchers.

Infrastructure varies widely: No two organizations are developing, hosting, or deploying generative AI in the same way. When everyone is using globally distributed facilities, different energy sources, and operational infrastructure, there's no such thing as a consistent measurement and reporting method.

Optimization hasn't been part of the equation: Because generative AI and its expanding demands are comparatively new, providers and users are only now turning their attention to optimizing the way models are created and operated.



¹<https://www.capgemini.com/insights/research-library/sustainable-gen-ai/>

²<https://www.capgemini.com/insights/research-library/sustainable-gen-ai/>

Charting a course for more sustainable AI operations

Almost a third of organizations are now incorporating sustainability measures into their generative AI projects.

But the United Nations University's Institute for Water, Environment and Health recently warned that

targeting individual components of AI's environmental footprint won't be enough to mitigate its impact. For example, using low-carbon electricity doesn't account for the amount of water required for cooling the facility.

It takes a holistic view to really make the difference.

Capgemini's methodology helps organizations take control of the resource risks and costs of AI across three key phases, uncovering the hidden layers of the value chain.

1 Anticipate

We assess your AI portfolio and strategy using our AI impact calculator and proven methodology to:

- Measure the footprint of your current AI infrastructure and portfolio
- Identify resource and cost hotspots
- Design a monitoring strategy that matches your deployments
- Create a single source of truth for the relevant data

2 Monitor

We deploy a Sustainable AI Cockpit with our Eco AI API to:

- Give you active control over current and future resource use
- Manage both AI products and AI capabilities within products
- Prioritize the most business critical and highest-risk applications
- Define and measure clear mitigation actions for the short, medium and long term
- Optimize and right-size models, architectures, and runtime efficiency
- Track trends, thresholds, and deviations across AI workloads

3 Embed

We set up a sustainable AI control tower with our library of technical, operational, and governance levers and learning resources, to:

- Assign ownership and responsibility
- Train end users to adopt responsible AI practices
- Equip developers with sustainable design and engineering techniques
- Empower stakeholders to make balanced decisions about performance and sustainability
- Make sustainability a default consideration of the whole AI lifecycle
- Review, challenge, and approve AI initiatives through established governance tools
- Engage partners and providers to increase transparency
- Build a system of active steering and continuous improvement



Working with Capgemini opens up a world of benchmarks and proven methodologies, helping organizations:

- Develop the fundamental data foundations, governance structures, and capabilities to run AI sustainably
- Build robust business cases, evaluate potential partners, and mitigate the footprint of active AI deployments
- Match real business needs to right-sized technologies, whether it's machine learning, more traditional AI models, or generative AI

Three pathways to more sustainable AI



80-90% impact from hardware | <10% of resources spent on training

Becoming the first AI-driven pharmaceutical company

How do you accelerate the development of AI-powered systems while maintaining tight climate commitments? This was the challenge for a large French pharmaceutical group.

Aiming for carbon neutrality by 2030 and net zero emissions by 2045, the group's AI ambitions needed firm guidelines. Capgemini used its Life Cycle Assessment methodology to build a multi-step calculator that considers the impact of factors including greenhouse gas emissions, resource depletion, water, and electricity consumption.

Our assessment found that up to 90% of the group's potential environmental impact was related to compute components, such as GPUs. Depending on the approach the team took, the footprint of its AI portfolio could become 10 times larger or reduce by 30% by 2030. With environmental cost-benefit scenarios in place, leaders can find the best proactive practices to implement more sustainable AI.



150k+ agents | 4B tokens | 45MWh electricity

Measuring and reducing the impact of agents

This major tech company offers an agentic AI platform that powers more than 150,000 autonomous agents in workflows for over 100,000 companies. The LLMs behind this service drive significant energy, carbon, and water usage and the associated financial costs. With a commitment to reducing its environmental footprint, this company needed precise measurement and an action plan.

In 10 weeks, Capgemini worked alongside the company to quantify the true impact of the platform, which uses 4 billion tokens each month. We identified hotspots in the architecture and processes that used the most resources, including 45MWh in monthly electricity demand.

Using this data, we created a robust environmental evaluation framework. Those insights combine with a user-friendly calculator and model optimization tools to help the team continuously monitor and manage the platform's impact.



15k agents = <0.01% GHG emissions | 1B agents = 60% GHG emissions

Exponential growth compounds resource risks

Many of the organizations we support are highly ambitious in their AI plans. But unlike some other technologies where costs reduce at scale, AI's demands only grow.

One of our clients had 15,000 agents in production in 2025. Leaders wanted to scale this to 1 billion in 2026, dramatically increasing AI's operational role across the organization. However, they knew that a jump of that scale would also increase the associated emissions and costs, so they brought in Capgemini to analyze their plans.

We found that the existing agents accounted for less than 0.01% of the company's greenhouse gas emissions in 2025. And if the team went ahead with its plans to expand, that would increase to 60% the following year.

Solve the AI paradox with Capgemini

Capgemini brings together recognized leadership in data, AI, and sustainability to deliver a more practical, responsible, and environmentally conscious approach to AI at scale.

Our combination of proven benchmarks and tailored frameworks helps leaders build transparency, auditability, and control into their generative AI

initiatives, so they can make informed, balanced decisions about everything from vendor partnerships to end user policies.

Explore how our methodology can help you unravel the AI paradox of impact and potential for your organization.

About the Authors



Philippe Cordier
Vice President & Global Chief AI Scientist
Capgemini Invent



Maik Schwalm
Global Sustainability Lead
Cloud Infrastructure Services
Capgemini



Franco Amalfi
Global Sustainability AI Partner Ecosystem Lead
Capgemini



Martin Chauvin
Manager Consultant
Data Scientist
Capgemini

For more details, contact:

Infra.global@capgemini.com

About Capgemini

Capgemini is the business transformation partner for enterprises in the age of AI. We help organisations imagine and build an intelligent, sustainable future, combining AI, technology and human ingenuity to transform how they operate, innovate and grow. With unique end-to-end capabilities spanning strategy, technology, engineering and intelligent operations, we bring together deep industry expertise and market-leading capabilities in AI, cloud and data to turn ambition into measurable business outcomes at scale. Supported by a robust ecosystem of partners and nearly 60 years of expertise, Capgemini is a responsible and diverse global organization of over 410,000 team members in more than 50 countries. The Group reported 2025 revenues of €22.5 billion.

Make it real.

www.capgemini.com

