


AI & Sustainability

Artificial intelligence is becoming foundational to how HubSpot builds products, runs internal systems, and supports customers. As we expand its use across the business, we are also evaluating the environmental impact of that growth, particularly in the absence of agreed industry standards or clear emissions data.

Our emerging approach to responsible innovation

AI-related emissions are challenging to measure. The majority of models are hosted externally, and providers offer limited visibility into infrastructure and energy consumption. While current data gaps remain, we're taking steps where we can to improve efficiency, strengthen internal guidance, and bring AI into closer alignment with our science-based carbon reduction targets.

Key Focus Areas



1. Infrastructure efficiency

For the cloud infrastructure we have control over, we aim to prioritize efficiency where appropriate. This includes:

- Running workloads during off-peak hours where possible.
- Selecting a third-party data center provider that operates on 100% renewable energy, and actively invests in efficient measures.

Where AI systems are hosted by external vendors, we have limited visibility into their infrastructure. These emissions are currently estimated as part of our Scope 3 footprint.

2. Model optimization

We've developed internal guidance to help teams choose the smallest effective model for the task at hand. These decisions reduce energy use, limit emissions, and improve performance across latency and cost.

3. Strategic hosting

We are exploring options to bring key models in-house. While most AI capabilities today rely on third-party APIs, internal hosting would provide more control over infrastructure choices, measurement, and long-term cost. This evaluation is ongoing.

4. Hardware

We're evaluating the use of more energy-efficient accelerators, such as GPUs and TPUs, for internal workloads. Where possible, we also partner with infrastructure providers that prioritize energy-efficient hardware.

5. Measurement

At present, we rely on spend-based Scope 3 estimates to account for emissions associated with vendor-hosted AI. These estimates are limited but reflect the best available methodology. As vendor reporting improves and new tools become available, we plan to expand our internal tracking and update our disclosures accordingly.

We're also working to better track AI-related energy use within our own infrastructure, starting with model training and high-volume inference workloads.

What comes next

Because we primarily rely on third-party cloud services and AI vendors, most of our AI-related emissions are currently indirect (Scope 3). This means we have limited control over the infrastructure itself, but we can influence outcomes through vendor selection, product design, and internal engineering choices.

Our strategy focuses on two parallel paths. First, we are optimizing areas where we have direct control: how models are selected, deployed, and scaled internally. Second, we are encouraging transparency and efficiency from vendors, recognizing that our long-term goals depend on collaboration across the AI ecosystem.

This will remain an active area of work as we continue to evaluate impact, refine controls, and improve data availability.