Is AI Helping or Hurting the Climate?

Assessment, Compliance, and Opportunity

Jean-Philippe EHRET

27/05/2026

resilio

MARS-S2L (Methane Alert and Response System)

Methane leaks detection system **Machine Learning** techniques and **Computer Vision**.

Verified, permanent mitigations

In Algeria only: prevented >27,000 tons of methane emissions annually, which is equivalent to 500,000 standard passenger vehicles' yearly emissions.

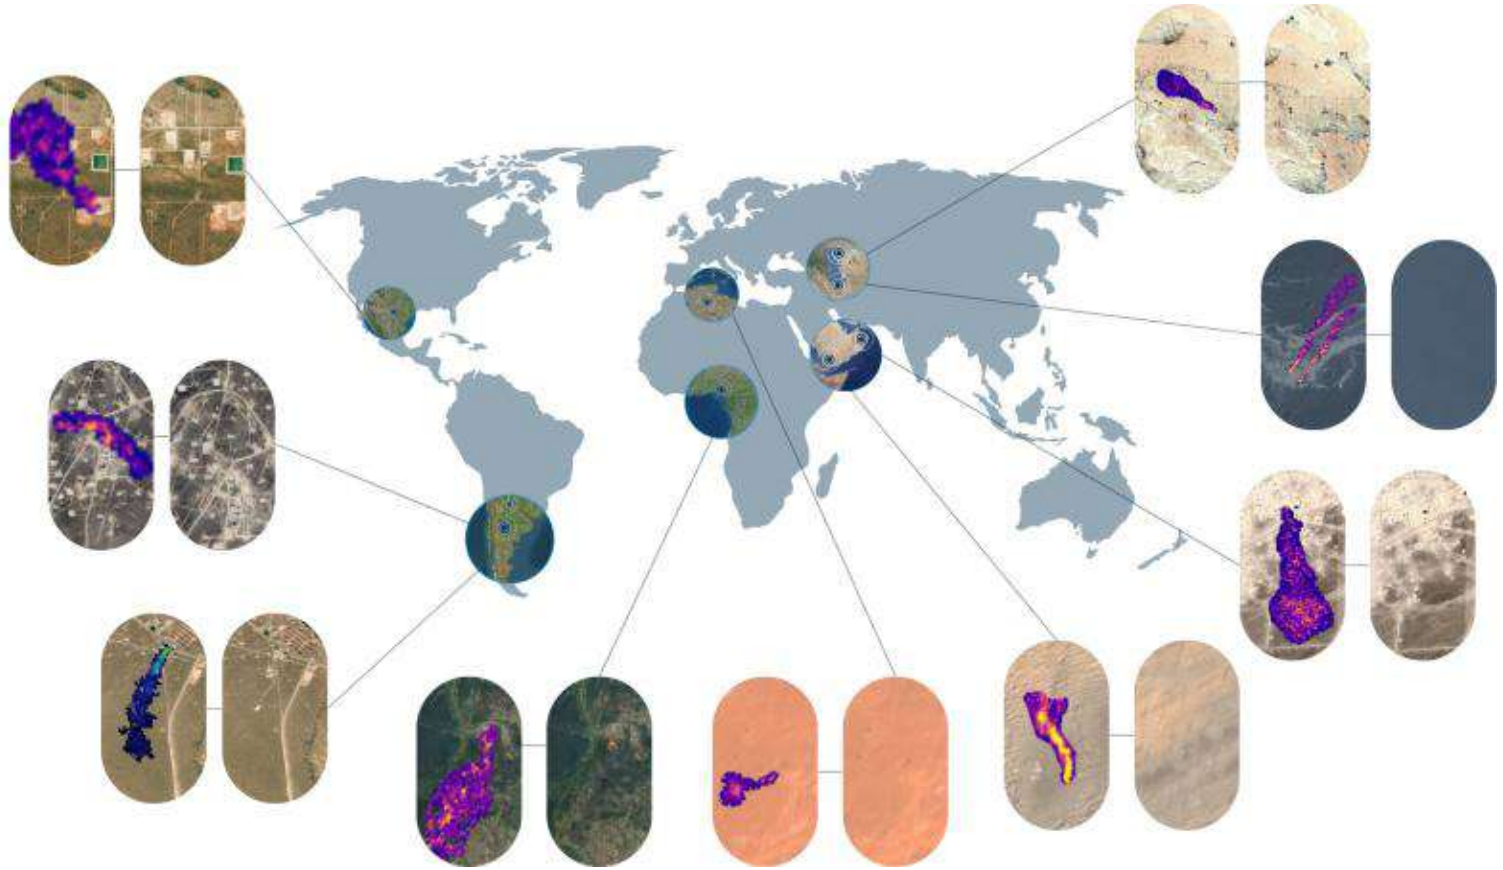

AI for Green/Handprint - A New Hope

AI for Green/Handprint - A New Hope?

A study of 154 AI climate claims found:
- 74% lack peer-reviewed evidence
- 36% cite no evidence at all

Source: Joshi, K. — "The AI Climate Hoax", Feb. 2026 / Beyond Fossil Fuels & Climate Action Against Disinformation

Meta identified 135 CO_2-absorbing materials using AI. Independent scientists found none worked — and some were completely fabricated.

Google's long-running claim that AI could mitigate 5–10% of global greenhouse gas emissions by 2030 based on a third-party blogpost containing unverified information.

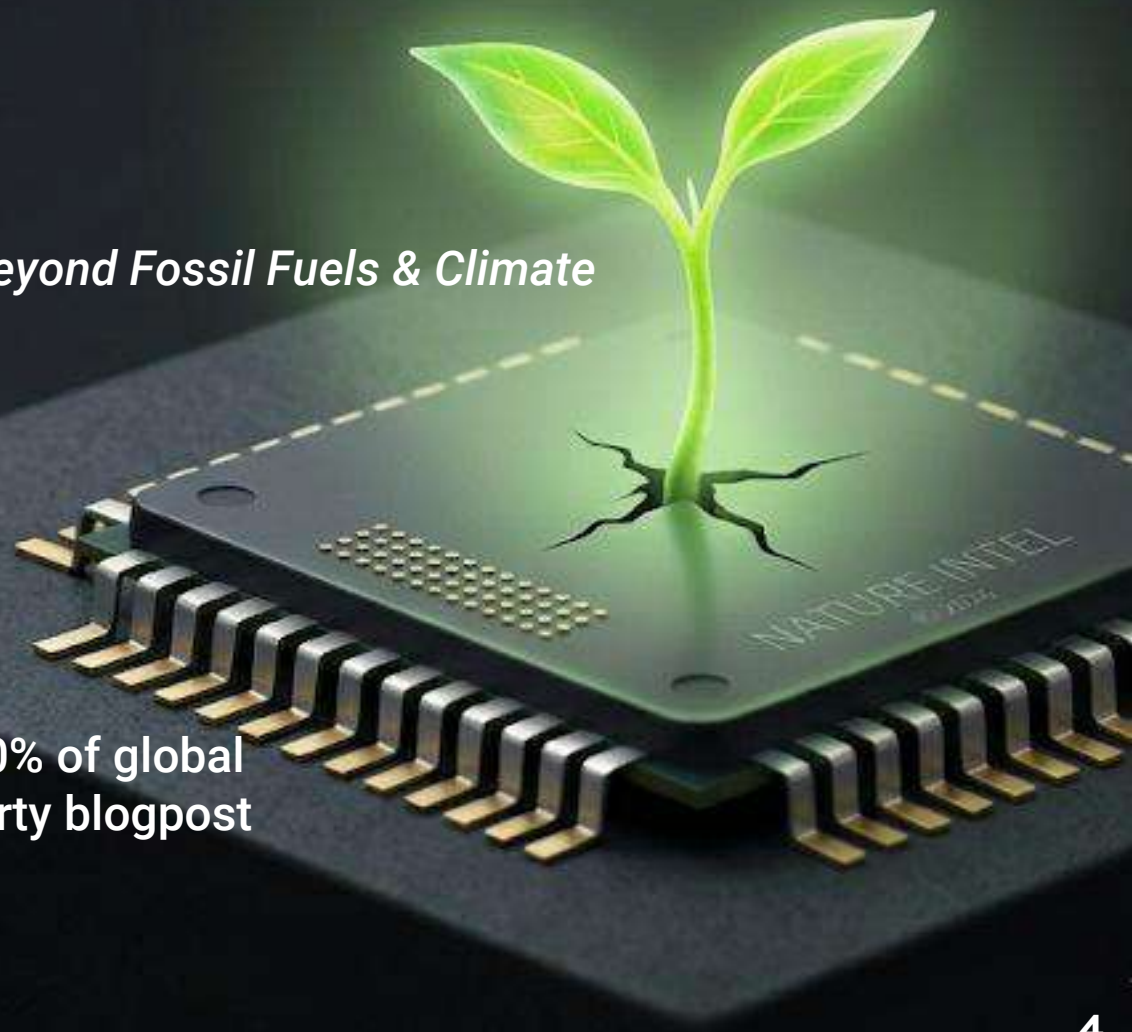

Not All AIs are the same

Traditional / Applied AI

Distribution of Claims

150 out of 154 total claims

97% of climate benefit narratives rely on older technologies like **predictive modeling** and **computer vision**

Generative AI

13x
Greater Energy Consumption

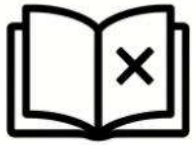

0

Zero climate claims are verified
(No consumer generative system has proven material emissions reduction)

Gigatonne-Scale Reductions are Unfounded

Evidence for massive AI climate benefits is weak, while evidence of data center harm is strong.

---> … At no point did this analysis uncover examples where consumer generative systems such as ChatGPT, Gemini, or Copilot led to a verifiable and substantial level of emissions reductions.

---> … much of the projected energy consumption of AI will stem specifically from “generative AI”, rather than more traditional forms of machine learning

The AI Climate Hoax: Behind the Curtain of How Big Tech Greenwashes Impacts
Ketan Joshi - February 2026

Environmental Impacts of LLMs

Tech companies successfully lobbied to have data centre energy and water consumption classified as *commercially sensitive* under EU law — making independent measurement nearly impossible.

Source: Investigate Europe / The Guardian, April 2026

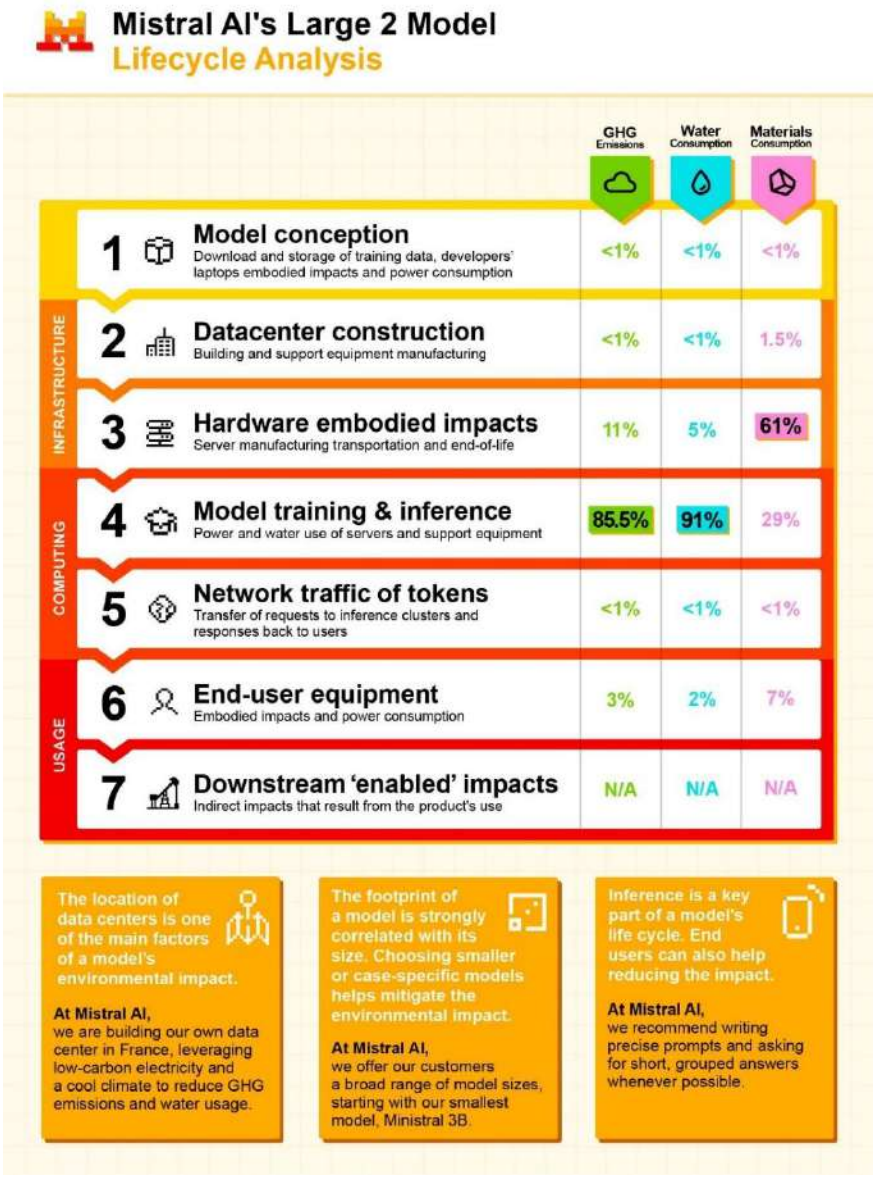

https://mistral.ai/news/our-contribution-to-a-global-environmental-standard-for-ai

Environmental Impacts of AI Datacentres

Land, Energy and Water are the hot topics.

Global AI power capacity is now comparable to peak power usage of New York State

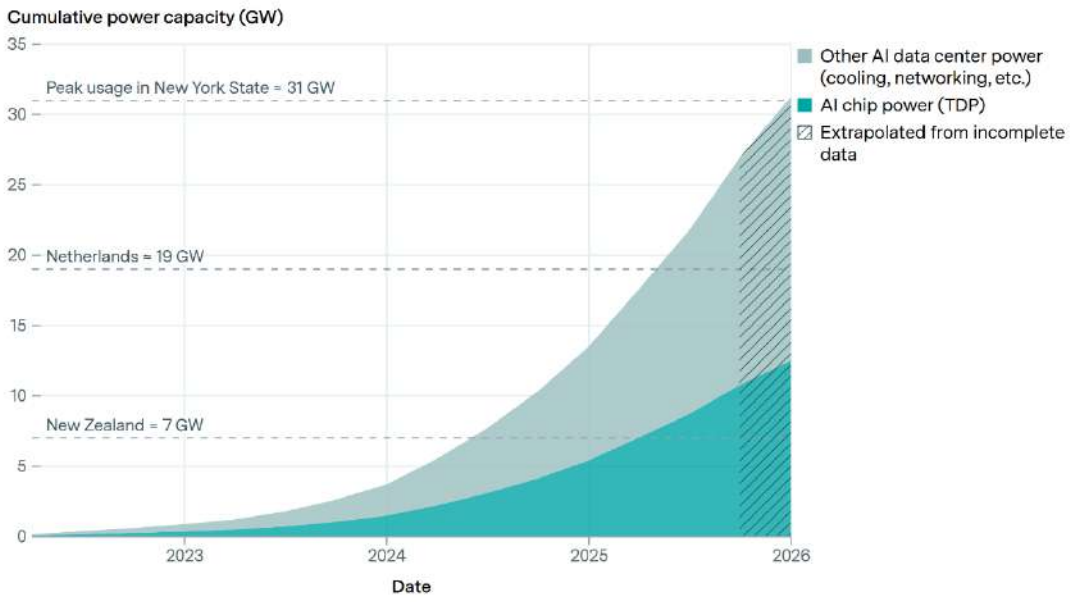

EPOCH AI | CC-BY epoch.ai

Global AI computing capacity is doubling every 7 months

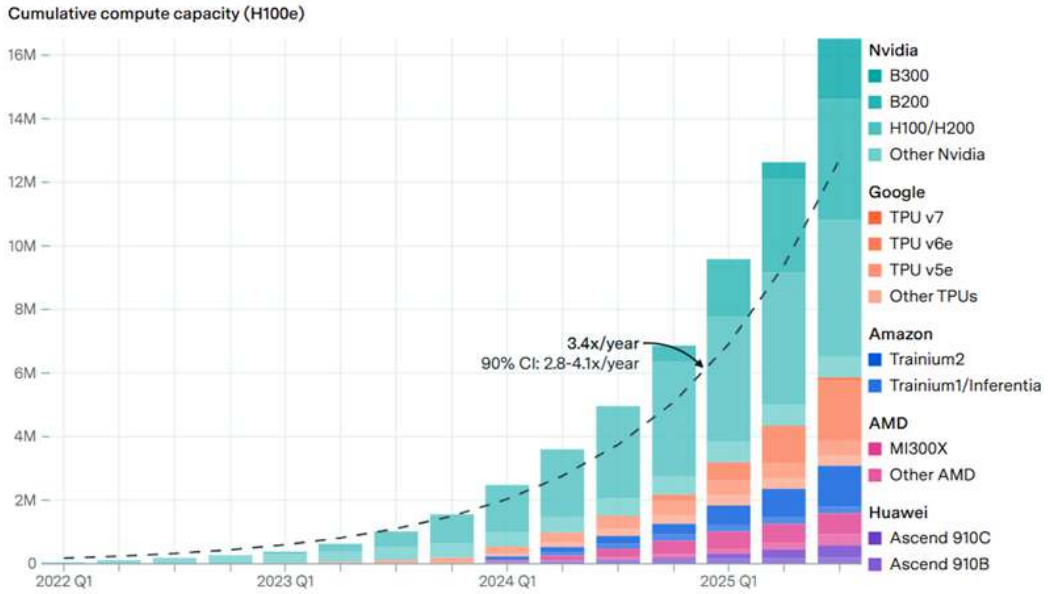

EPOCH AI | CC-BY

AI vs. Societal Essentials

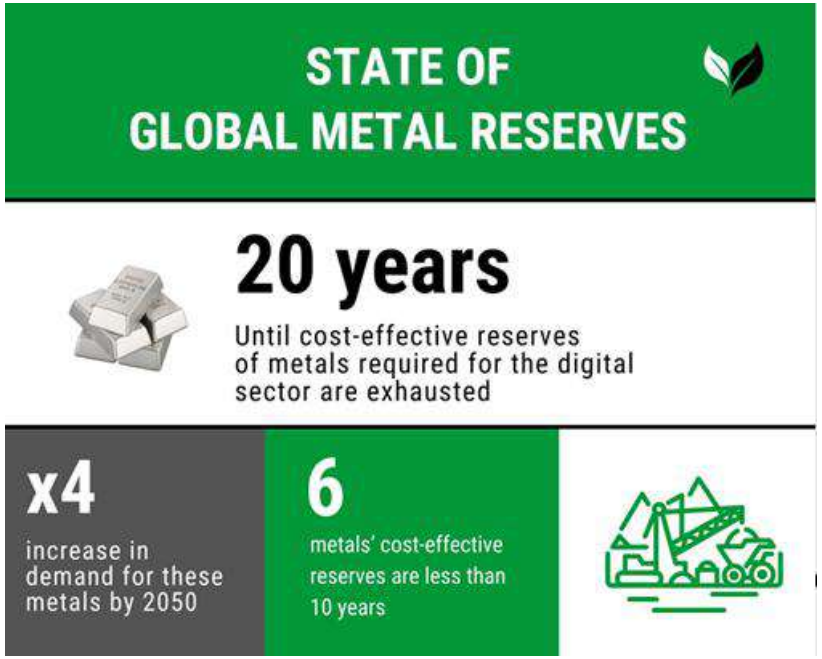

How long can Generative AI remain a Mass Market product?

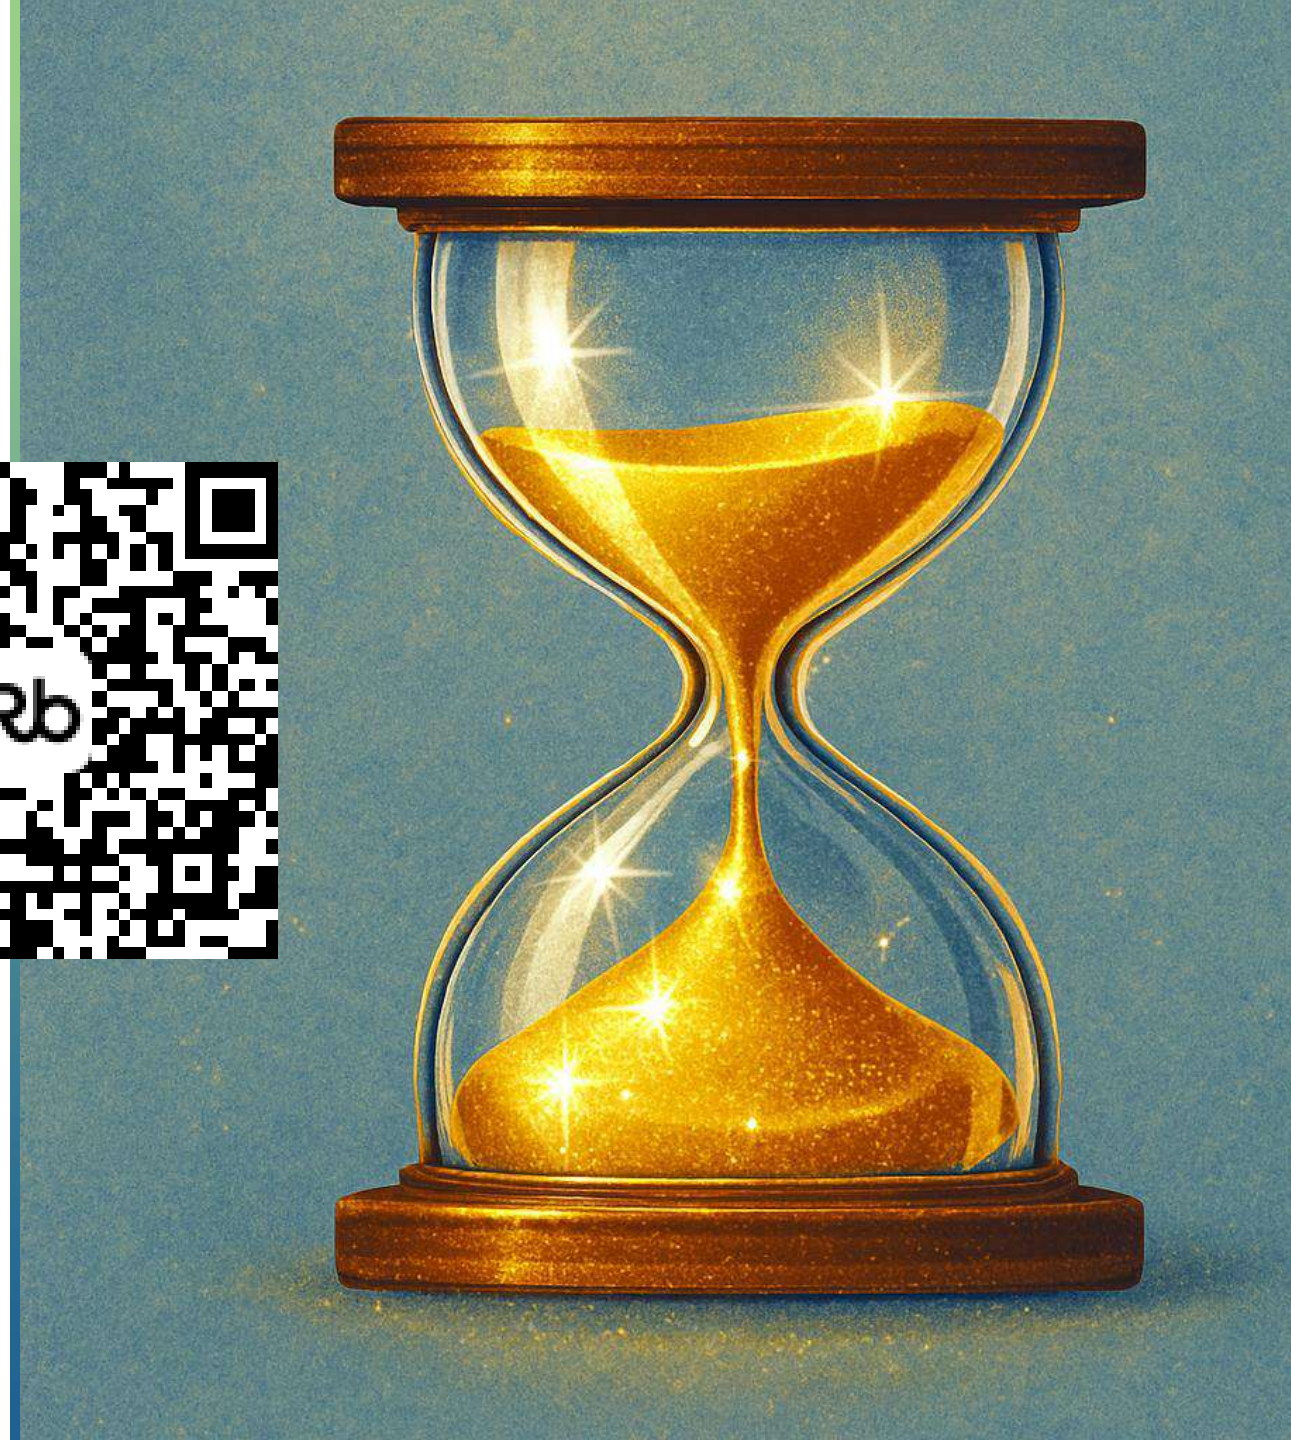

The Optimisation Case/Enabled Emissions

Efficiency ≠ Sustainability as it comes with **indirect impact** and **rebound effect**

Optimisation is available to everyone. Including the fossil fuel industry.

"Microsoft's AI contracts with ExxonMobil and Chevron enabled an estimated **57 million additional tons of CO_2** — more than 3× Microsoft's own annual emissions." — *Global Witness, 2026*

AI Governance – Environmental Sustainability

- Regulatory & disclosure risk regarding environmental impact
- Reputational & greenwashing risk
- Cost & operational efficiency risk
- …

AI Forum New Zealand
Te Kāhui Atamai Iahiko o Aotearoa

AI Governance – Environmental Sustainability

Climate-related disclosures (CRD)
AI can materially affect reported energy use and supplier impacts.

Environmental certifications
Programmes such as Toitū Envirocare, ISO 14001, and B Corp require identification of material environmental impacts, digital and AI services cannot be excluded.

Procurement and suppliers
Cloud and AI vendors increasingly influence sustainability performance and disclosure quality

Where organisations operate overseas,
extraterritorial requirements (e.g., EU AI Act) may also apply.

Materiality assessment: make impact visible

IS AI MATERIAL?

Proportionate response:

AI READY › VISIBILITY
AI SMART › MEASUREMENT
AI TRUSTED › OPTIMISATION

ITU-T L.1801

• New Standard · ITU-T Study Group 5 · Approved 06 February 2026

AI finally has an environmental report card.

ITU-T L.1801 is the world's first standard for measuring how much AI systems actually cost the planet — from GPU mining to your last query. This is what it means, layer by layer.

4	5	3
Life cycle stages every AI must account for	Environmental impact categories defined	Orders of effect — including ripple effects on the world

TYPE 01
Expert Systems
Medical diagnosis, fraud detection, energy optimisation — rule-based, human-defined logic
Energy
Data
Hardware: CPU only

TYPE 02
Machine Learning
Email filtering, product recommendations — learns from statistical patterns in data
Energy
Data
Hardware: CPU + GPU

TYPE 03
Deep Learning
Image recognition, CNNs, RNNs — multiple layers of neural networks
Energy
Data
Hardware: CPU + GPU

TYPE 04 Highest Impact
Generative AI
LLMs, image/video generation, multi-modal models — generating new content at scale
Energy
Data
Hardware: Significant GPUs and/or TPUs

→ **Key insight:** The standard explicitly classifies GenAI as **"comparatively high"** in both energy and hardware demands — a level above all other AI types. When people debate "AI's environmental impact," they are mostly talking about Type 04.

Materiality assessment: make impact visible

Extend you IT Inventory to include environmental drivers

When organisations start looking, they find:
- **far more AI than they thought**
- **only a small share that's environmentally material**
- the **ones material far more material than they thought**.

INVENTORY DATA: The four questions main questions
- **Where/When does it run?** — cloud provider, region, what's the grid like
- **How much does it run?** — volume, frequency of inference but also frequency of retraining (including R&D)
- **What kind of AI is it?** — output format (Text, Image, Video), **Reasoning, Agentic**
- **Who's exposed?** — which suppliers, which disclosures (CRD), which rules (EU AI Act, certifications)

At this stage, coarse is fine. The goal is to separate the few material systems from the many that aren't.

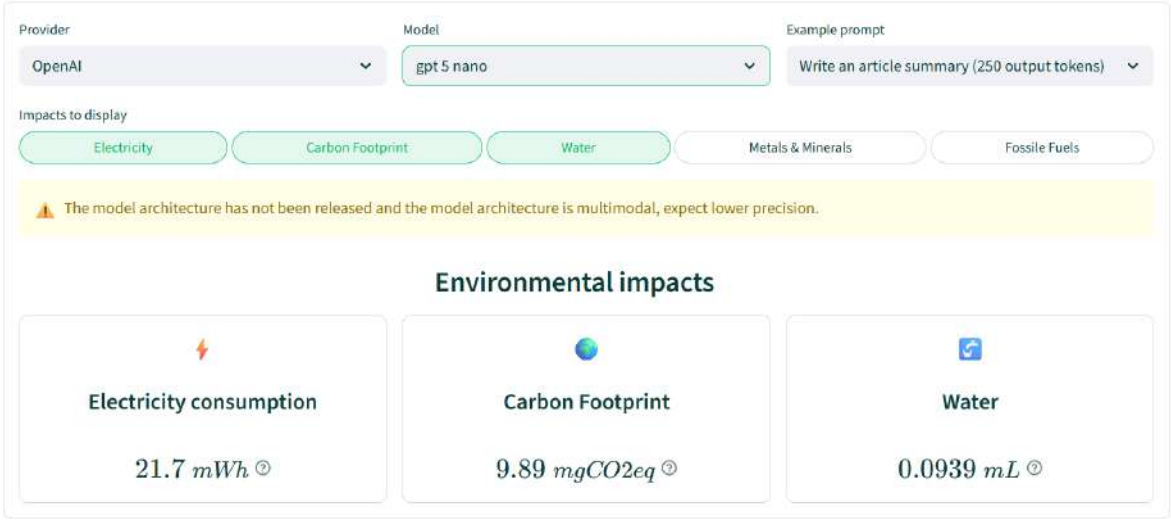

EcoLogits

Estimate and track the environmental footprint of GenAI models at inference.

Made with by CodeCarbon non-profit.

AgenticAI, a major shift in Sustainability

Recursive, parallel, switches on storage, APIs, jobs, autonomously and sometimes for hours.

A single agentic task can use around *a **1000x*** more compute than a standard AI question.

Impacts entire IT infrastructures, within and outside the organisation.

LCA & GreenOps

AI providers burned **tens of billions in 2025** to offer free access to **hundreds of millions of users**.

Price has been deliberately **decoupled from cost**.

Measure tokens, not seats or licenses. And monitor the impact of token consumption on your entire IT infrastructure, not only the AI compute.

Spend-based emission factors don't work here (they never worked for IT anyway).

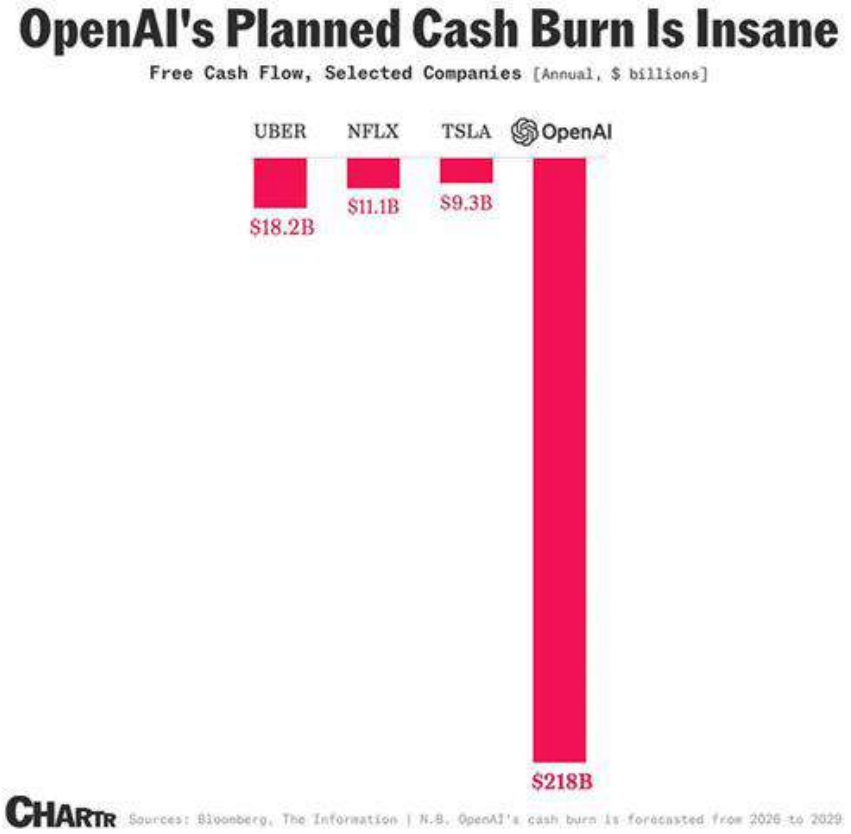

Scope 3 can represent over **90%** of an org. total carbon footprint, especially in the services sector.

When measured appropriately, **AI exacerbates Digital Emissions** across categories **1**, **2**, **11**, **12** and **15**.

Mitigation Principles

- Right-size models
- Choose lower-carbon regions/time
- Implement eco-design principles for efficiency
- Reduce unnecessary inference & retraining
- Extend hardware lifetime & support circular procurement
- Green Prompting

Right-sized models

"Stop using a sledgehammer to crack a nut."
Defaulting to the largest general-purpose model wastes **20–30× the energy of a fit-for-purpose smaller one** — for no extra value.

(Luccioni, Hugging Face)

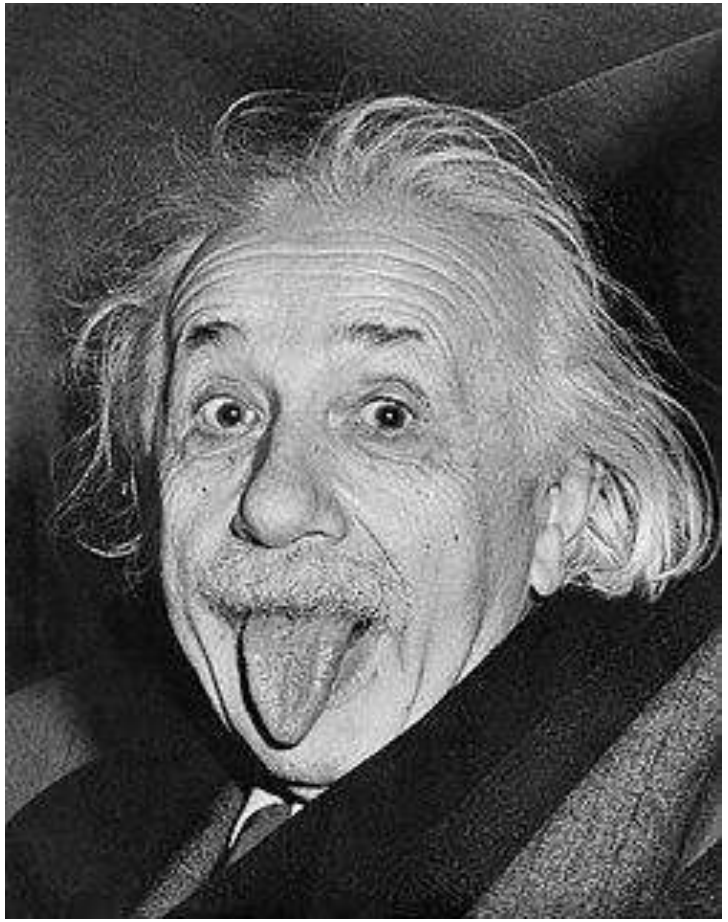

Region/Time

75% of the US data centers currently being built are set to be powered by natural gas.

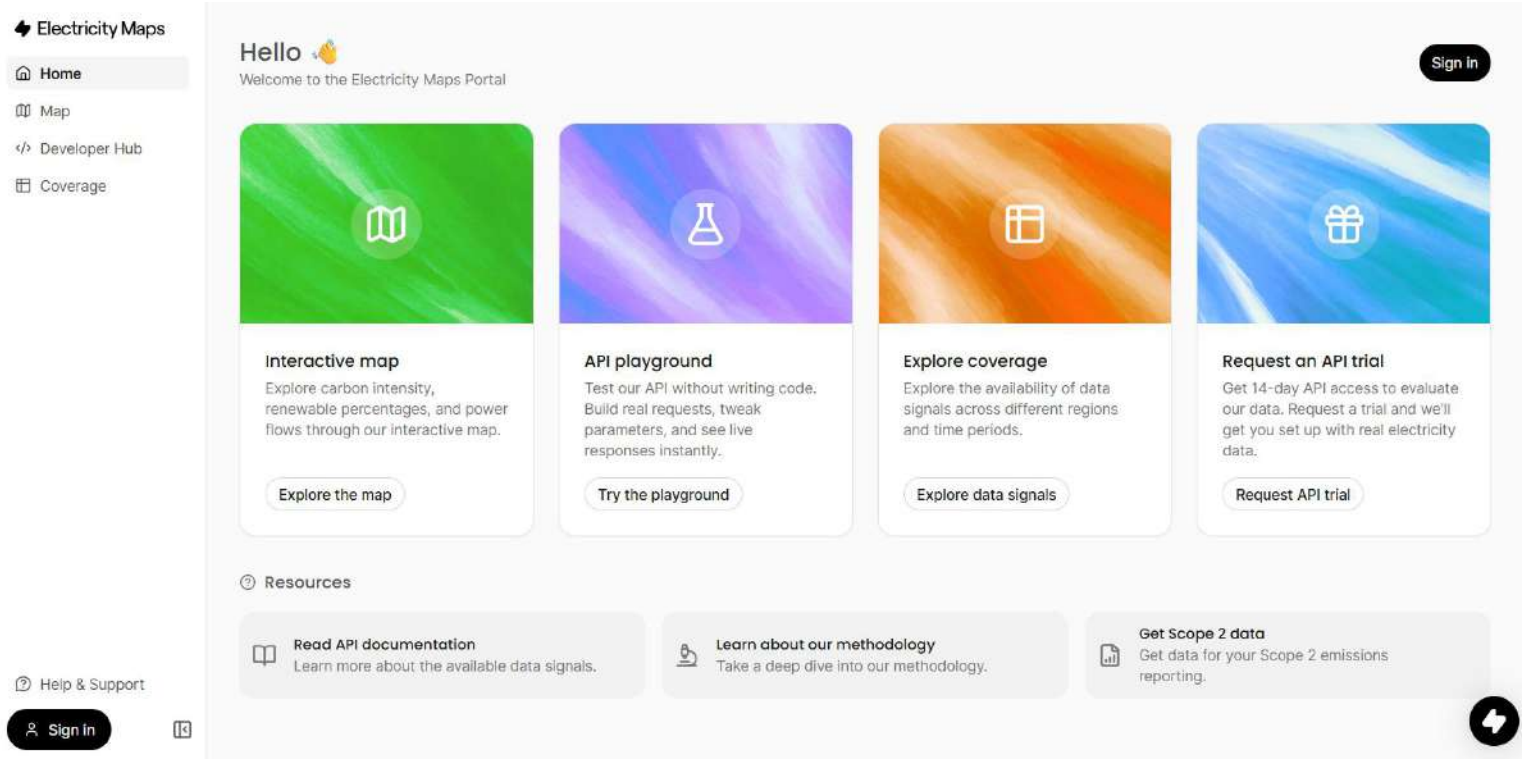

Easy: Run your AI at the right place, at the right time.

Advanced: Adopt Load Shifting/Grid Aware computing.

Environmentally Friendly Chatbot - Euria

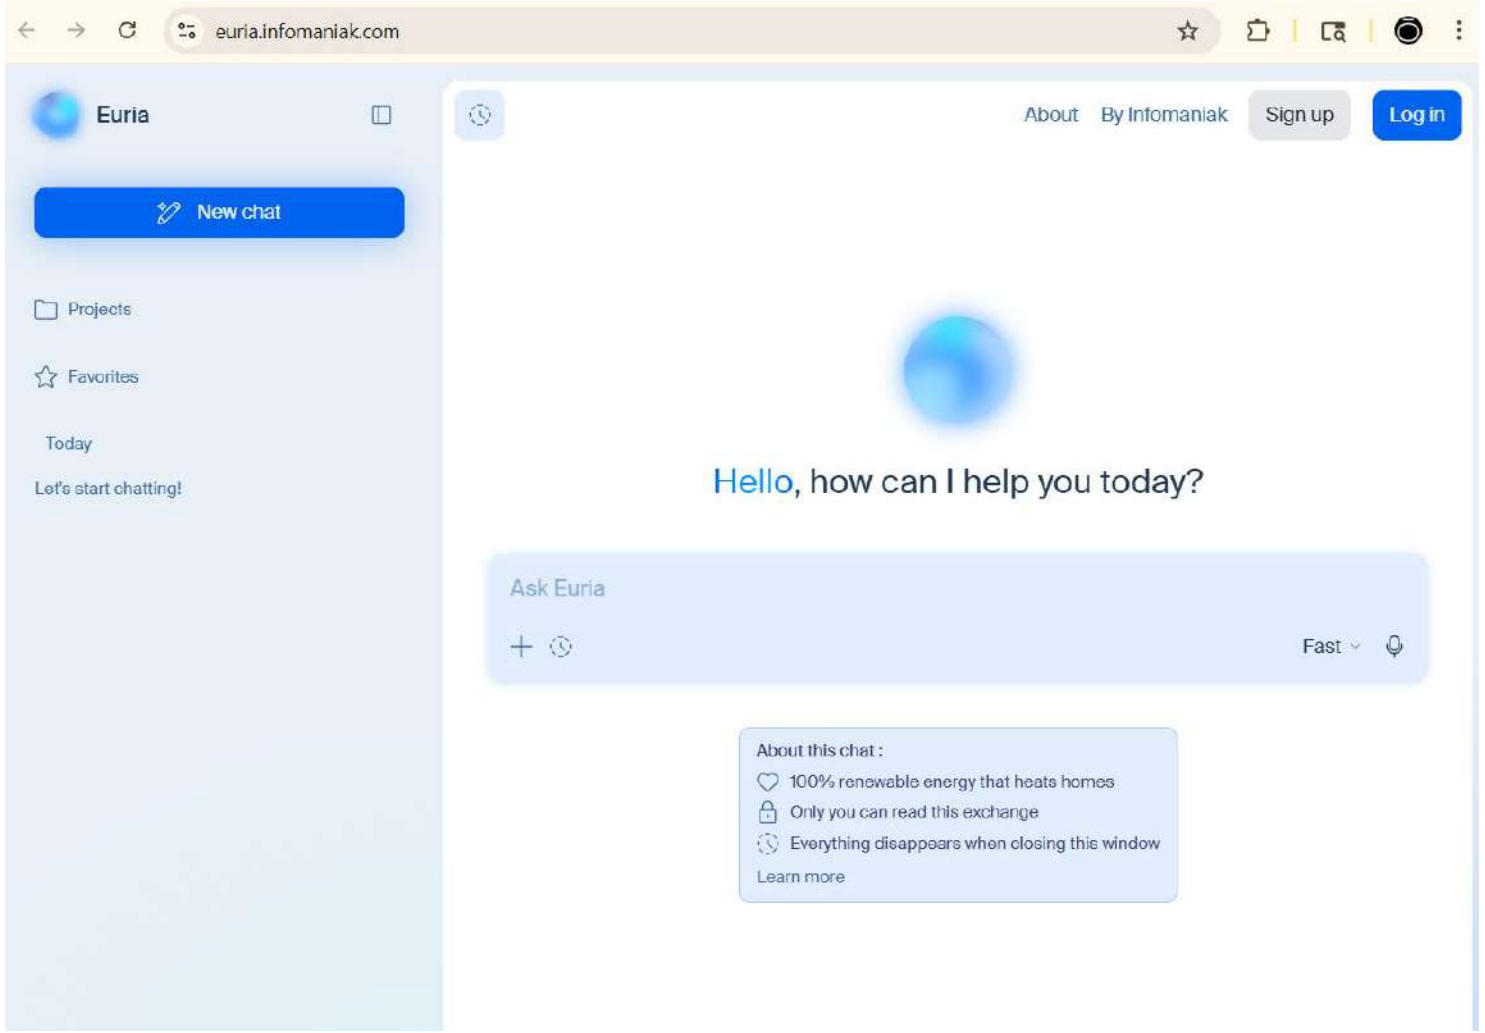

The Greenest Prompt? Keep Conversations Short

*Because transformer attention scales **quadratically** with context length, turn 20 doesn't cost 20× turn 1, it costs closer to 400×.*

Summarize and restart to reduce **energy** and **cost** (compact/handoff)

> **give me a conversation handoff summary designed to restart a fresh session with zero loss of context. Structure it as:**
>
> 1. **My goal / what I'm trying to achieve** (2–3 sentences)
> 2. **Key decisions made and conclusions reached** (bullet points)
> 3. **What we tried and ruled out** (brief)
> 4. **Open questions or next steps** (bullet points)
> 5. **Critical context or constraints I should carry forward** (anything the AI needs to know about me, my project, my constraints)
>
> Write it as something I can paste at the top of a new conversation as the very first message, so the new session picks up exactly where this one left off. Be concise — this is a brief, not a transcript."**

***Adjust to your context/model*

The Greenest Prompt? LLMs don’t need literature

README Contributing MIT license

caveman

why use many token when few do trick

stars 61k last commit last tuesday license MIT

Before/After • Install • What You Get • Benchmarks • Full inst

A Claude Code skill/plugin (also Codex, Gemini, Cursor, Windsurf, Cline, Copilot, 30+ caveman — cuts ~75% of output tokens, keeps full technical accuracy. Brain still big.

Normal Claude

> "Sure! I'd be happy to help you with that. The issue you're experiencing is most likely caused by your authentication middleware not properly validating the token expiry. Let me take a look and suggest a fix."

Caveman Claude

> "Bug in auth middleware. Token expiry check use `<` not `<=`. Fix:"

Same fix. 75% less word. Brain still big.

TOKENS SAVED	▓▓▓▓▓▓	75%
TECHNICAL ACCURACY	▓▓▓▓▓▓	100%
SPEED INCREASE	▓▓▓▓▓▓	~3x
VIBES	▓▓▓▓▓▓	OOG

Eco-Design

Eco-design your software: Delete non-essentials

Allow a budget for it.
It is worth it!

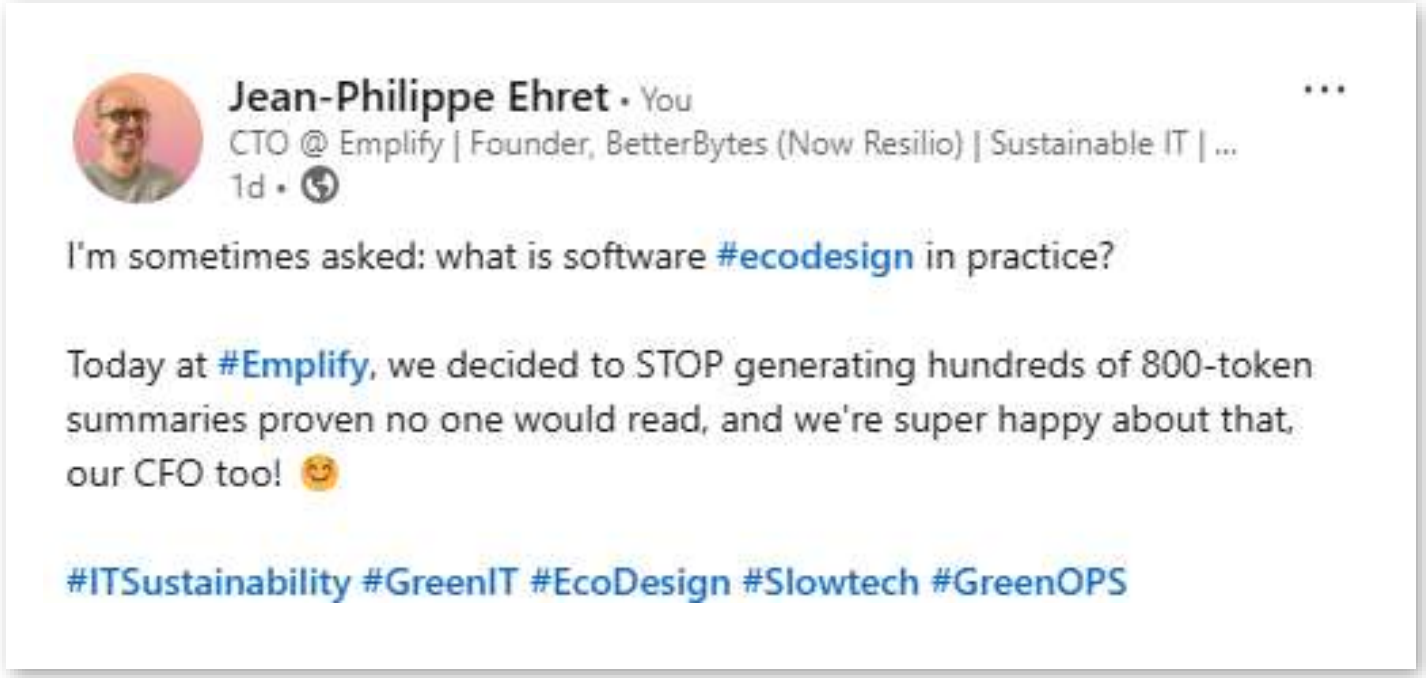

Take-aways

- Traditional/Applied AI proven positive for climate; not consumer GenAI.
- Generative AI impacts energy, carbon, water, and raw materials — impact resilience and sovereignty.
- All organisations must embed environmental sustainability in their AI Governance.
- Materiality assessment: always. Measurement: likely when CRD, certifications, extraterritorial rules, or agentic AI are present.
- AI's footprint extends to the entire IT infrastructure, agentic AI amplify impact significantly
- Right-size models, choose lower-carbon regions/times, apply eco-design — and remember optimisation can enable others' emissions too.
- Green Prompting: Not a silver bullet, but a few best practice worth knowing

Together, let's build a more sustainable digital future

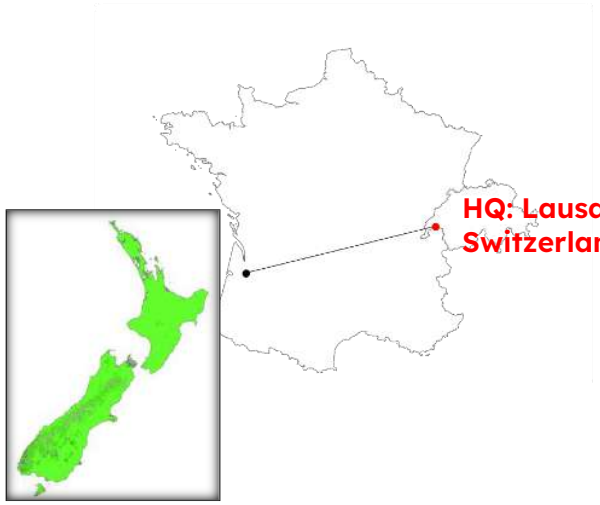

Certified B Corporation

Resilio SA is a **Swiss** based company helping organisations define and implement **Sustainable IT** strategies.
Its **NZ** based subsidiary is based in Hawke's Bay and offering services in APAC.

Questions?

resilio