The Unseen Cost of Intelligence: AI as the Fastest-Growing Sustainability Challenge — and How to Account for It

A White Paper by Cetus AI *June 2026*

Executive Summary

Artificial intelligence has become the defining productivity technology of the 2020s. Enterprises across every sector are deploying large language models (LLMs) for customer service, software development, legal review, financial analysis, and strategic decision-making. Yet the environmental cost of this deployment is almost entirely invisible in corporate sustainability reporting. Every token generated by an AI model consumes electricity. That electricity produces carbon. And under the GHG Protocol's Corporate Value Chain Standard and Australia's mandatory AASB S2 climate disclosure regime, those emissions belong on the balance sheet.

This white paper makes three arguments. First, AI-driven energy consumption is growing faster than any other sector of the global economy and is on track to become a material emissions source for most large enterprises within the next two to three years. Second, the mechanism by which AI emissions accumulate — token by token, API call by API call — is structurally invisible to existing sustainability reporting frameworks, creating a compliance gap that regulators are beginning to close. Third, Songlines Control provides the first purpose-built platform for capturing, attributing, and disclosing AI token emissions in a format that satisfies GHG Protocol Scope 3 Category 1 requirements and the AASB S2 mandatory disclosure standard.

1. The Scale of the Problem: AI's Accelerating Energy Footprint

1.1 Data Centres Are the Fastest-Growing Energy Consumer on Earth

Global electricity consumption from data centres reached approximately 415 terawatt hours (TWh) in 2024, representing around 1.5% of total global electricity demand. That figure has grown at 12% per year over the past five years — but the pace is accelerating sharply. The International Energy Agency (IEA) projects that data centre electricity consumption will **double to approximately 945 TWh by 2030**, growing at around 15% per year — more than four times the rate of growth of all other electricity-consuming sectors combined. [1]

The primary driver of this acceleration is artificial intelligence. Accelerated servers — the GPU and TPU clusters that power AI inference and training — are projected to grow their electricity consumption at **30% per year** in the IEA's base case, compared to 9% per year for conventional servers. By 2030, AI-related workloads are expected to account for 35–50% of all data centre power consumption, up from approximately 5–15% today. [2]

To put this in context: the electricity consumed by global data centres in 2024 was roughly equivalent to the annual electricity consumption of Thailand. By 2030, on current trajectories, it will approach the consumption of Japan.

1.2 The Token Is the Unit of Emissions

Unlike a factory or a fleet of vehicles, AI emissions do not come from a single identifiable source. They accumulate invisibly, one API call at a time. Each time an enterprise system sends a prompt to an LLM and receives a response, a measurable quantity of electricity is consumed and a measurable quantity of carbon is emitted.

Research published in 2025 provides the most granular picture to date of per-token emissions across the major model families. [3] The figures vary substantially by model architecture, hardware generation, and the carbon intensity of the data centre's electricity grid:

Model Scale	Hardware	Estimated CO_2e per Output Token
288B+ parameters (e.g., GPT-o3)	NVIDIA H100 @ FP16	~30 mg CO_2e
70B parameters	NVIDIA A100 @ FP16	~15 mg CO_2e
70B parameters	NVIDIA H100 @ FP8	~7.5 mg CO_2e
13–27B parameters	NVIDIA A100 @ FP16	~3 mg CO_2e
2B parameters	NVIDIA A100 @ FP16	~0.5 mg CO_2e

Source: DitchCarbon synthesis of "From Words to Watts" benchmarks and vendor performance data, 2025[3]

A single ChatGPT request emits approximately **4.32 g CO_2e** on average — roughly twenty times the emissions of a Google search (≈ 0.2 g CO_2e). [3] A 1,000-token image-enhanced request has been measured at 8.3 g CO_2e — equivalent to charging a smartphone ten times. [3]

These figures seem trivial in isolation. They are not trivial at enterprise scale.

1.3 The Enterprise Emissions Gap

Consider a mid-sized professional services firm with 500 employees, each making an average of 20 AI-assisted interactions per working day across tools such as Microsoft Copilot, GitHub Copilot, Claude, and internal LLM-powered applications. At 4.32 g CO_2e per interaction, that firm generates approximately **10.8 tonnes of CO_2e per year** from AI inference alone — before accounting for the substantially higher emissions of agentic workflows, code generation, and multimodal tasks that consume far more tokens per interaction.

For a large enterprise with tens of thousands of employees and AI embedded across customer-facing, operational, and back-office functions, the figure scales to hundreds or thousands of tonnes per year. A 2026 analysis estimated that AI-enabled coding tools alone are generating approximately **250,000 tonnes of carbon emissions annually** across the industry. [4]

Critically, **none of this appears in most corporate sustainability reports**. The emissions are real, they are attributable, and they are growing — but the accounting frameworks to capture them have lagged behind the technology.

2. The Regulatory Imperative: Why This Is Now a Compliance Issue

2.1 The GHG Protocol and Scope 3 Category 1

The GHG Protocol's Corporate Value Chain (Scope 3) Standard, which underpins virtually every major corporate emissions reporting framework globally, divides a company's value chain emissions into 15 categories. **Category 1 — Purchased Goods and Services** covers all upstream emissions from the production of goods and services that a company acquires. [5]

AI inference services — API subscriptions, enterprise chat tools, SaaS products with embedded AI features — fall unambiguously within Scope 3 Category 1. When a company pays OpenAI, Anthropic, Google, Microsoft, or AWS for AI API access, it is purchasing a service whose production generates greenhouse gas emissions. Those emissions are the company's Scope 3 Category 1 liability.

A peer-reviewed paper published in June 2026 confirmed this classification and proposed a four-tier methodology for including AI inference in corporate GHG inventories: [6]

> *"AI inference services — API subscriptions, enterprise chat tools, and SaaS products with embedded AI features — fall unambiguously within Scope 3 Category 1 under the Corporate Sustainability Reporting Directive (CSRD), which requires disclosure for fiscal years starting January 2024. Yet no standardised methodology exists for including them in corporate GHG inventories. Current practice either omits the category entirely or applies a generic economic input-output (EEIO) factor calibrated to the ICT sector as a whole, overestimating AI inference emissions by 10–40x relative to physically derived alternatives."*

The paper's key finding is that the compliance challenge is **methodological rather than magnitude-driven**: the emissions themselves are relatively small for most organisations, but the absence of a rigorous, auditable methodology means that companies cannot demonstrate compliance even when they want to. This is precisely the gap that Songlines Control is designed to close.

2.2 Australia's Mandatory AASB S2 Disclosure Regime

Australia has moved faster than most jurisdictions to make climate disclosure mandatory. The Australian Accounting Standards Board's **AASB S2 — Climate-related Financial Disclosures** standard, released in September 2024, requires entities to disclose information about climate-related risks and opportunities that could reasonably be expected to affect their cash flows, access to finance, and cost of capital. [7]

Mandatory reporting commenced on 1 January 2025, with a phased rollout by entity size:

Group	Revenue Threshold	Mandatory From
Group 1	> $500M revenue	FY2025 reports (lodged mid-2026)
Group 2	$50M$–500M revenue	FY2026 reports
Group 3	< $50M revenue	FY2027 reports

Source: AASB S2 (Sep 2024); EY Mandatory Climate-related Financial Disclosures Summary[7] [8]

Scope 3 emissions are required from the **second reporting year** for each group, meaning that Group 1 entities must include Scope 3 disclosures in their FY2026 reports — lodged in late 2026 or early 2027. For organisations that have not yet begun capturing AI token emissions data, the window to establish a baseline is closing rapidly.

AASB S2 aligns closely with the International Sustainability Standards Board's IFRS S2, meaning that compliance with AASB S2 substantially satisfies reporting obligations under the CSRD (for entities with European operations), the SEC's climate disclosure rules (for US-listed entities), and the UK's Streamlined Energy and Carbon Reporting (SECR) framework.

2.3 The Reporting Gap in Practice

Despite these obligations, the current state of AI emissions reporting is characterised by three systemic failures:

Omission. The majority of corporate GHG inventories do not include AI inference emissions at all. The category is either unknown to sustainability teams or treated as immaterial without any quantitative basis for that assessment.

Misclassification. Where AI emissions are reported, they are frequently misclassified — absorbed into general IT energy consumption (Scope 2) or estimated using generic ICT sector input-output factors that overstate actual emissions by 10–40 times. [6]

Attribution failure. Even where emissions are captured at the organisational level, they cannot be attributed to specific business units, workflows, models, or providers — making it impossible to identify reduction opportunities or demonstrate the effectiveness of mitigation measures.

---

3. The Structural Invisibility of Token Emissions

3.1 Why Existing Tools Cannot Solve This Problem

Traditional sustainability reporting tools are designed around physical energy consumption — electricity meters, fuel receipts, travel records, and supplier spend data. They are not designed to capture the granular, real-time, per-request data that AI emissions accounting requires.

The fundamental unit of AI emissions is the **token** — a sub-word fragment of text, roughly 0.75 words on average. A single API call to a large language model may consume thousands of tokens. The emissions from that call depend on the specific model, the number of input and output tokens, the hardware generation on which the model runs, and the carbon intensity of the electricity grid at the data centre where the inference occurs.

None of this information is captured by a corporate expense management system, an ERP, or a traditional carbon accounting platform. The data exists — it flows through API logs, telemetry systems, and usage dashboards — but it is not connected to emissions accounting infrastructure.

3.2 The Compounding Effect of Agentic AI

The emissions challenge is being compounded by the rapid adoption of **agentic AI** — autonomous systems that chain multiple LLM calls together to complete complex tasks. A single agentic workflow — such as a software development agent that writes, tests, and deploys code — may consume hundreds of thousands of tokens across dozens of sequential API calls. The emissions from a single agentic task can exceed those of thousands of standard queries.

Research published in 2026 estimated that AI coding agents alone are generating approximately 250,000 tonnes of CO_2e annually across the industry. [4] As agentic AI moves from developer tools into enterprise operations — automating procurement, customer service, financial analysis, and strategic planning — the emissions from these workflows will become a material line item in corporate GHG inventories.

3.3 The Water Dimension

Carbon is not the only environmental cost of AI inference. Data centres consume substantial quantities of water for cooling. Research has identified a **water-carbon trade-off** that existing ESG tools do not surface: data centres in regions with low-carbon electricity grids (such as Sweden, which relies heavily on hydropower) often have the highest water footprint per unit of compute, with direct implications for data centre location strategy and supply chain sustainability assessments. [6]

A widely cited estimate suggests that ChatGPT consumes approximately 500 millilitres of water for every 5–50 prompts processed. [9] At enterprise scale, this water consumption is a material environmental impact that belongs in sustainability disclosures alongside carbon emissions.

---

4. Songlines Control: Purpose-Built for AI Emissions Accountability

4.1 What Songlines Control Does

Songlines Control is an AI governance and sustainability platform purpose-built for enterprises that use large language models at scale. It captures token-level telemetry from every AI API call made across an organisation's technology stack, attributes emissions to specific models, providers, business units, and workflows, and generates audit-ready disclosure reports aligned with GHG Protocol Scope 3 Category 1 requirements and AASB S2.

The platform operates through a lightweight telemetry layer that integrates with existing AI infrastructure — whether that is direct API calls to OpenAI, Anthropic, Google, or AWS, or AI features embedded in enterprise SaaS platforms. No changes to existing workflows are required. Emissions data is captured automatically, in real time, and attributed at the granularity required for both operational decision-making and regulatory disclosure.

4.2 The Emissions Calculation Methodology

Songlines Control calculates emissions using a **token-based physical estimation** methodology — the highest-precision tier in the four-tier framework proposed by Llopis (2026). [6] For each API call, the platform records:

- The model name and provider
- The number of input tokens and output tokens
- The timestamp of the request
- The cost of the request (where available)

Emissions are calculated using the formula:

> ***CO_2e (g) = total_tokens ÷ 1,000,000 × model_emissions_factor***

Emissions factors are sourced from peer-reviewed research, including Lannelongue et al. (arXiv:2505.09598, 2025) and Devera AI's published benchmarks, and are updated as new data becomes available. The Australian grid intensity figure of 0.51 kg CO_2e/kWh is sourced from the Department of Climate Change, Energy, the Environment and Water's National Greenhouse Accounts (2024). [10]

An uncertainty range of ±30% is applied to all estimates, reflecting variability in data centre location, renewable energy mix, and hardware generation — consistent with the GHG Protocol's guidance on Scope 3 estimation uncertainty.

4.3 The AASB S2 Report

Songlines Control generates a structured AASB S2 AI Emissions Report that covers all mandatory disclosure elements:

Report Section	Content
Reporting Entity	Entity name, ABN, financial year, reporting period
Emissions Classification	Standard (AASB S2), GHG Protocol category, activity description, providers
Emissions Summary	Total CO_2e in grams, kilograms, and tonnes; total tokens; total requests; models used
Calculation Methodology	Formula, emissions factor sources, grid intensity, uncertainty range, boundary
Mandatory Disclosure Applicability	Group classification, reporting timeline, Scope 3 phase-in
Model Breakdown	Per-model token consumption, request volume, CO_2e contribution, percentage share

The report is generated on demand, covers any specified financial year, and is formatted for direct inclusion in annual sustainability reports or submission to regulators. It is designed to satisfy the disclosure requirements of AASB S2, IFRS S2, CSRD, and the GHG Protocol Corporate Value Chain Standard simultaneously.

4.4 Beyond Compliance: Operational Intelligence

Songlines Control is not only a compliance tool. The same token-level telemetry that powers regulatory disclosure also enables organisations to:

Identify high-emissions workflows. By attributing emissions to specific workflows, business units, and use cases, organisations can identify where AI usage is generating disproportionate environmental impact and prioritise optimisation efforts accordingly.

Optimise model selection. The platform surfaces the emissions cost of different model choices in real time. Switching from a 70B parameter model to a 13B parameter model for tasks that do not require frontier capability can reduce emissions by 80% or more, with no loss of output quality for many use cases.

Track emissions intensity over time. As AI usage scales, organisations can monitor whether their emissions intensity (CO_2e per unit of AI output) is improving or deteriorating — a key metric for demonstrating progress against net-zero commitments.

Benchmark against peers. Aggregated, anonymised data from across the Songlines Control customer base enables organisations to benchmark their AI emissions intensity against industry peers — a capability that will become increasingly important as investors and regulators demand comparative disclosure.

---

5. The Business Case for Acting Now

5.1 Regulatory Risk

For Group 1 entities in Australia — those with revenue above $500 million — the first Scope 3 disclosures under AASB S2 are due in late 2026 or early 2027. Organisations that have not established a baseline for AI token emissions by mid-2026 will face one of two outcomes: either they omit the category from their disclosure (creating a material gap that auditors and regulators will increasingly scrutinise) or they rely on generic EEIO estimates that may overstate their actual emissions by 10–40 times, creating reputational and legal exposure in the opposite direction. [6]

The cost of establishing a rigorous, auditable methodology now is a fraction of the cost of retrofitting one under regulatory pressure.

5.2 Investor and Stakeholder Expectations

Institutional investors, including Australia's major superannuation funds, are increasingly requiring portfolio companies to disclose Scope 3 emissions as part of their climate risk assessments. A 2026 analysis estimated that inaccurate Scope 3 data creates up to **$500 billion in hidden corporate risk** globally, through regulatory penalties, lost green premiums, and investor flight. [11]

AI emissions are a small but rapidly growing component of Scope 3 for most large enterprises. Organisations that can demonstrate rigorous, granular AI emissions accounting will be better positioned to satisfy investor due diligence requirements and maintain access to sustainability-linked financing.

5.3 The First-Mover Advantage

The organisations that establish AI emissions accounting infrastructure now will be the ones that can demonstrate year-on-year improvement in emissions intensity as AI usage scales. This is a significant competitive advantage in an environment where sustainability performance is increasingly a factor in procurement decisions, talent attraction, and brand reputation.

Songlines Control enables organisations to move from a position of ignorance — where AI emissions are unknown and unaccounted for — to a position of leadership, where AI usage is governed, optimised, and transparently disclosed.

---

6. Conclusion

The environmental cost of artificial intelligence is not a distant or hypothetical concern. It is accumulating right now, in every enterprise that uses AI-powered tools, at a rate that is growing faster than any other category of corporate energy consumption. The regulatory frameworks that require disclosure of these emissions are already in force in Australia, and the window to establish a compliant baseline is measured in months, not years.

The challenge is not the magnitude of the emissions — for most organisations, AI token emissions are currently a small fraction of total Scope 3 — but the absence of the infrastructure to measure, attribute, and disclose them accurately. Existing sustainability reporting tools were not designed for this task. Generic estimation methods produce results that are simultaneously unreliable and unauditable.

Songlines Control closes this gap. By capturing token-level telemetry from every AI API call, applying peer-reviewed emissions factors, and generating structured AASB S2-compliant disclosure reports, it gives organisations the data they need to meet their regulatory obligations, satisfy investor expectations, and make informed decisions about how they use AI.

The token is the unit of AI value. It is also the unit of AI emissions. Organisations that understand and account for both will be better positioned to govern AI responsibly — and to demonstrate that governance to the stakeholders who increasingly demand it.

References

Songlines is an AI governance and sustainability platform built by Cetus AI for enterprises that use large language models at scale. For more information, contact Cetus AI at hello@cetusai.com.au.

1. International Energy Agency. *Energy Demand from AI — Energy and AI*. IEA, 2025. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai
2. Carbon Brief. *AI: Five Charts That Put Data-Centre Energy Use — and Emissions — Into Context*. Carbon Brief, 2025. https://www.carbonbrief.org/ai-five-charts-that-put-data-centre-energy-use-and-emissions-into-context/
3. DitchCarbon. *The Real Carbon Cost of an AI Token*. DitchCarbon Blog, April 2025. https://ditchcarbon.com/blog/llm-carbon-emissions
4. CNaught. *AI Coding Agents Are Emitting 250,000 Tonnes of Carbon Emissions Each Year*. CNaught Blog, May 2026. https://www.cnaught.com/blog/ai-coding-agents-are-emitting-250-000-tonnes-of-carbon-emissions-each-year-heres-how-we-got-there
5. GHG Protocol. *Technical Guidance for Calculating Scope 3 Emissions — Chapter 1: Purchased Goods and Services*. World Resources Institute, 2013. https://ghgprotocol.org/sites/default/files/2022-12/Chapter1.pdf
6. Llopis, G. *Accounting for AI Inference in Corporate GHG Inventories: A Four-Tier Methodology for Scope 3 Category 1 Reporting*. arXiv:2606.10660, June 2026. https://arxiv.org/abs/2606.10660
7. Australian Accounting Standards Board. *AASB S2 — Climate-related Financial Disclosures*. AASB, September 2024. https://standards.aasb.gov.au/aasb-s2-sep-2024
8. EY Australia. *Mandatory Climate-related Financial Disclosures: What You Need to Know*. EY, 2025. https://www.ey.com/content/dam/ey-unified-site/ey-com/en-au/insights/sustainability/documents/ey-mandatory-climate-related-financial-disclosures-update-and-summary-july.pdf

9. Li, P. et al. *Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models*. arXiv:2304.03271, 2023.
10. Department of Climate Change, Energy, the Environment and Water. *National Greenhouse Accounts Factors 2024*. Australian Government, 2024.
11. CO2 AI. *The $500B Price Tag of Inaccurate Scope 3 Data*. CO2 AI Insights, February 2026. https://co2ai.com/insights/the-hidden-cost-of-inaccurate-scope-3-data-compliance-risk-lost-premium-opportunities-and-investor-skepticism