# Artificial Intelligence, Sustainability, and the Discipline to Govern Both

By Dr. Joel Carboni | Founder & President, GPM | Director, PMI-GPM

> *AI is not a software abstraction. It is a physical system. Before we talk about what it can do for sustainability, we have to be willing to talk honestly about what it does to the planet just by running.*

Artificial intelligence is everywhere in professional conversation right now. Every conference agenda has a slot for it. Every leadership team is under pressure to have a position on it. That pressure is real, and the technology is moving fast enough that staying quiet feels like falling behind.

But here is what keeps getting left out of that conversation: AI depends on data centers, electricity grids, water for cooling, semiconductors, and mineral supply chains that span the globe. That starting point is not pessimism. It is basic intellectual honesty. And it is the foundation for a conversation that actually moves somewhere useful.

PMI's research in Sustainability in the Age of AI makes the enterprise case plainly. Organizations that integrate AI into their sustainability strategy outperform those that treat the two as separate programs, and the report names integration as the imperative. Its sharper finding is that execution, not intent, separates leaders from the rest, with only a quarter of organizations reporting that they execute their sustainability initiatives successfully. That is the gap this paper addresses. Integration is declared in the boardroom, and it is decided in the project. The business case, the vendor selection, the delivery governance, and the post-implementation review are where AI either earns its footprint or fails to. PMI already supplies the instruments for that work. The PMI® GPM® P5 Standard for Sustainability in Project Management governs sustainability impact at the point of delivery, and the Trustworthy AI Framework, part of the CPMAI™ Training program governs how AI behaves once it is running. This paper is about requiring their use.

## THE FOOTPRINT IS REAL AND IT IS GROWING

Energy demand from AI is rising, not stabilizing. The International Energy Agency projected that data centers could consume more than 950 terawatt-hours of electricity annually by 2030, roughly equivalent to Japan's entire national consumption. Water used for cooling and chip manufacturing runs into billions of liters per year at individual facilities. It is worth distinguishing between water consumed and water returned: much of the water used in cooling processes is returned to the source, often cleaner than when it entered. But water that is evaporated is lost entirely, and that volume is material at scale. The materials embedded in the hardware carry extraction and supply chain costs that do not appear on any vendor's pricing sheet.

Reporting on these impacts varies significantly across the industry. Some major cloud providers have published renewable energy commitments and achieved meaningful progress toward carbon goals. The more important question is what percentage of energy consumption was drawn from actual renewable generation before credits were purchased, and whether those ratios have been maintained as AI demand has accelerated. Credits and commitments are not the same as real-time renewable supply. That distinction matters when organizations are making infrastructure choices based on sustainability claims.

The question is not whether the footprint exists. It does. The question is whether the organizations deploying AI are measuring it, disclosing it, and including it in their value calculations. Most are not.

*"AI is the most powerful capability humanity has ever created, while sustainability remains the most important outcome humanity must achieve. The challenge and opportunity is ensuring the first accelerates the second."*

**Lloyd Skinner**
CEO, Greyfly.ai

## WHY THIS MATTERS BEYOND THE FOOTPRINT

Governing AI well is not just about limiting harm. It is about preserving the capacity to use AI where it genuinely counts.

Climate modeling at the scale required to inform infrastructure decisions depends on computational power that no other approach provides. Early warning systems for water stress, flooding, and food supply disruption rely on pattern recognition across datasets too large for conventional analysis. Reducing emissions across industrial supply chains requires optimizing thousands of interdependent variables simultaneously: logistics routes, energy sources, material substitutions, and emissions accounting in real time. Detecting ecosystem degradation across planetary boundaries requires continuous monitoring of satellite, sensor, and biological data at a resolution no human team can process manually. These are active applications. They represent AI being used for work that cannot be done any other way.

Beyond solving discrete problems, AI is being used to redesign how organizations operate. Manufacturers are using it to identify material substitutions that reduce waste across product lines. Logistics companies are optimizing routes and load factors in ways that cut emissions structurally, not just operationally. Financial institutions are using it to model the long-term viability of low-carbon investment portfolios that conventional risk models cannot evaluate at sufficient resolution. These are not sustainability reports being generated faster. They are business models being restructured.

That is precisely why undisciplined use is a problem that extends beyond any single organization. An organization that deploys AI for tasks that did not need it consumes grid capacity, water supply, and political tolerance for data center expansion that serious applications also depend on. Those resources are shared and finite. Treating AI as a default solution for every problem is not just wasteful. It draws down the infrastructure that could otherwise support the applications we actually need.

## THE STANDARD THAT IS CURRENTLY MISSING

The problem is not that organizations lack information about AI's environmental cost. That information is available. The problem is that no one in most approval chains is currently required to answer a straightforward question: does this deployment create enough value to justify what it consumes?

Sustainability teams and AI teams operate in separate lanes. Procurement moves faster than governance. The people making deployment decisions are rarely the ones who carry the downstream costs. Until the proportionality question is a required part of the approval process, the answer will always be assumed to be yes. That is a structural gap, and it produces predictable results.

Part of applying that standard is recognizing that AI solutions exist on a wide spectrum of resource intensity. A lightweight, purpose-built model that solves a defined problem consumes a fraction of the energy and infrastructure of a large general-purpose system. Choosing the most capable or most prominent tool when a simpler one would do is not a neutral decision. It is a cost that someone pays, even if it does not show up on the project budget. Right-sized, efficient AI is not a compromise. It is the professional standard.

*"If you thought AI was complex... the environmental impact is even more complex."*

*"There is quite a bit of data cherry picking as the demand for AI expands. We need to be thorough and thoughtful at the same time."*

**Ben Royce**

Global Chief Technology Officer, AKQA | Formerly Google Cloud AI Services

## GOVERNANCE HAS TO REACH THE PROJECT LEVEL

Policy frameworks and principles exist in reasonable supply. UNESCO established an ethical framework for AI. The OECD and the Global Partnership on Artificial Intelligence have published governance principles. ISO/IEC 42001, an international management system standard published in 2023, gives organizations a structured way to govern AI at the policy and controls level. These are real contributions. The problem is that none of them have produced measurable change in how most organizations make AI deployment decisions, because they all stop before they reach the person writing the business case, selecting the vendor, or approving the scope.

In every organization, decisions happen at the project level. Projects are where AI becomes operational, where procurement choices are made, where architecture is chosen, where implementation scope is defined, where risk is allocated, and where expected value gets tested against actual results. If sustainability impact is discussed only at the strategy level, it stays rhetorical. For it to be real, it has to show up in the business case, the vendor evaluation criteria, the delivery governance process, the benefits realization framework, and the post-implementation review.

That is the territory of project professionals. And it is where PMI and GPM have a specific, non-delegable responsibility.

## QUESTIONS EVERY PROJECT LEADER SHOULD BE ASKING

For project professionals, this is not abstract. It comes down to a set of questions that should be standard on any project where AI is being evaluated or deployed.

Is there a real problem being solved, or is this deployment driven primarily by market pressure? If the honest answer is the latter, that is a sufficient reason to pause. No footprint is justified by trend-following alone.

Is there a simpler, less resource-intensive approach that could produce the needed outcome? Proportionality means choosing the tool that fits the problem, not the one with the most capability or the largest marketing budget.

Can the vendor provide clear, verifiable information on the infrastructure implications of the system being proposed? If the answer is no, that belongs in the evaluation. Opacity on environmental cost is not a neutral data gap. It is a risk.

If AI is being used in support of sustainability goals, do the underlying data, assumptions, and governance mechanisms hold up to scrutiny? A model built on poor data or governed by opaque rules does not become trustworthy because it carries a sustainability label.

These are not anti-innovation questions. They are the same management discipline that project professionals apply to capital expenditure, regulatory risk, and vendor selection. AI as a tool deserves the same scrutiny as any other, not a carve-out.

## WHAT THIS MEANS FOR PROJECT MANAGERS WHO ARE NOT BUILDING AI

Most project managers are not deploying AI systems. They are using them. Tools like Claude, Microsoft Copilot, and ChatGPT have become part of daily work: drafting documents, summarizing meetings, generating options, reviewing plans. That use is not neutral, and it is not outside the scope of this conversation.

Every query sent to a large language model runs on data center infrastructure with a real energy and water footprint. A single query is small. But project management is a profession of millions of practitioners worldwide. When that community adopts AI-assisted tools as a routine part of daily work, the aggregate consumption becomes significant. It rarely appears in any sustainability report, which makes it one of the more consequential unmanaged impacts in the profession right now.

The P5 Standard's Responsible Technology element provides the starting point. It applies to any project that uses technology, not just projects that build it. It asks project teams to identify ethical, legal, and operational risks associated with deployed technologies, establish governance mechanisms for AI and automated decision systems, and evaluate the environmental impacts of digital infrastructure where material. That scope covers the project manager using Copilot to draft a charter just as much as the team building a machine learning pipeline.

In practice, that means three things. First, use these tools for work that genuinely benefits from them. A task where AI surfaces an insight or removes genuine tedium is a reasonable use. Generating output you could have produced in the same time, or that you will not actually use, is waste. The proportionality test applies at the task level, not just the enterprise level.

Second, when evaluating or recommending AI tools, push for vendor transparency on infrastructure. Data privacy and security are already standard evaluation criteria. Energy use, water consumption, and carbon reporting from the vendor should be part of that conversation too. If a vendor cannot provide that information, it is worth naming in the evaluation. The Responsible Technology element gives you the professional standard to cite when you raise it.

Third, know your organization's policy on AI tool use. If one does not exist, say so. The absence of a policy is not a free pass. It is a governance gap that sits inside every project that touches these tools, and surfacing it is part of doing the job well.

None of this requires becoming an environmental specialist. It requires the same habit of mind that good project managers already bring to scope, risk, and value: ask whether the work is justified, whether the decision is documented, and whether someone is accountable for the outcome. P5 already defines that as professional practice. AI tools are not exempt from it.

## THE CREDIBILITY GAP ORGANIZATIONS NEED TO CLOSE

Organizations cannot make strong public claims about sustainability and simultaneously treat AI as exempt from scrutiny. If sustainability reporting covers supply chains, operations, and capital projects, then AI deployment belongs in that scope too. Exempting it is not a technical decision. It is a choice, and stakeholders are increasingly aware of it.

The PMI-GPM P5 Standard for Sustainability in Project Management provides the structure to do this work. It covers people, planet, prosperity, process, and products, with assessment frameworks that function at the project delivery level. Unlike ESG reporting frameworks designed for corporate disclosure, or organizational management system standards like ISO/IEC 42001, which establishes how institutions govern AI at the policy and controls level, P5 works where decisions actually happen, inside the project. PMI's CPMAI certification extends this further into AI-specific territory, with the Trustworthy AI Framework addressing ethical use, responsible deployment, transparency, governance, and explainability across the full AI project lifecycle. Together, these two frameworks cover the ground that most organizations are currently leaving unmanaged: the sustainability impact

of AI at the point of delivery, and the governance of AI behavior once it is running. The project profession already has the tools. The question is whether organizations are willing to require their use.

Several jurisdictions are already moving on mandatory AI disclosure and impact assessment. Organizations that build governance into their project delivery frameworks now will be better positioned as those requirements take effect. Getting ahead of that is both more defensible and less expensive than retrofitting accountability after the fact.

The case for responsible AI governance does not rest on regulatory compliance alone. Organizations have an obligation to account for what they consume and what they create, including the costs imposed on people, communities, and ecosystems that had no say in the decision. That obligation does not disappear because a technology is commercially attractive.

## WHAT NEEDS TO HAPPEN NOW

The argument in this paper is not that AI should be used less. It is that the project profession has a specific responsibility that it is not currently fulfilling. Every AI deployment passes through a project. Every project has a business case, a governance structure, a benefits realization plan, and a post-implementation review. Those are exactly the points where sustainability impact should be measured, disclosed, and weighed against value created. The infrastructure to do this already exists. The P5 Standard and the CPMAI Trustworthy AI Framework are not proposals. They are operational tools available now. What is missing is the organizational requirement to use them.

For PMI and GPM, the practical agenda is this: make sustainability governance a required component of AI project delivery, not an optional add-on. That means updating governance standards, certification requirements, and delivery frameworks to reflect that AI projects carry environmental and social costs that have to be accounted for. It means giving project professionals the language, the tools, and the organizational backing to ask hard questions and to push back when the answers are not good enough. The profession has done this before with safety, with data privacy, with risk management. AI as a tool is not different. It is just newer.

The conversation has spent long enough on whether AI and sustainability are compatible. They can be, under the right conditions. The question worth asking now is who is responsible for creating those conditions, and whether the project profession is ready to step into that role.

Green Project Management.

**About the Author**

Dr. Joel Carboni is the Founder and President of GPM (Green Project Management) and Director at PMI-GPM. He holds a Ph.D. in Sustainable Development and Environment and is the architect of the P5 Standard for Sustainability in Project Management. He is a visiting professor at Skema Business School and a member of the Forbes Business Council