

AI Agents in Sustainability

Lessons from the Big Four's Industrial-Scale Deployment

Adam Wynne | adam@wynnetech.ai

March 23, 2026

Executive Summary

Sustainability teams are being asked to do more with less: more frameworks, more regulatory pressure, more data, while remaining one of the smallest functions in most organizations. AI agents offer a way out of that bind. But for most sustainability leaders, AI adoption still feels like a leap of faith.

It doesn't have to. The Big Four consulting firms have already taken that leap, at industrial scale, and the architecture they built for general business functions maps directly onto ESG work.

This whitepaper makes three contributions:

1. **Documents verified Big Four AI deployments** and extracts the patterns that matter for sustainability
2. **Maps those patterns to six ESG application categories**, showing where proven agent architectures apply
3. **Introduces the AI Maturity Model for Sustainability Practitioners** — the first maturity framework built specifically for ESG teams, assessing four dimensions across five levels

The core insight is simple: the tools already exist. They are running under different names in audit and tax. This whitepaper shows how to apply them to sustainability.

1. The Deployment Reality: What's Actually Happening

The shift from AI experimentation to production deployment happened fast. By late 2025, all four major professional services firms had named platforms, dedicated agent ecosystems, and published outcome metrics.

1.1 Verified Deployment Scale (Late 2025)

Firm	Production Agents	In Development	Platform
EY	150 specialized tax agents (supporting 80,000 professionals)[4]	~850 more[2]	EY.ai Agentic Platform
PwC	250+ platform agents + 25,000 client-deployed agents[1]	Growing	Agent OS
KPMG	50 agents[5]	~1,000[5]	Workbench
Deloitte	Not disclosed (470k employees with Claude AI)[6]	Unknown	Zora AI

PwC Agent OS explicitly lists ESG reporting as a core capability.[7] These aren't separate ESG tools. They're the same multi-model orchestration platforms used for audit and tax, applied to sustainability use cases.

1.2 Quantified Outcomes

PwC is the only Big Four firm to publish specific outcome metrics from production deployments:[7]

- 8x faster cycle times in targeted processes
- 30% cost reduction with human oversight maintained
- 70% reduction in manual compliance review time

These aren't projections. They're measured results from production systems that apply directly to ESG data collection and reporting workflows.

2. The Market Opportunity

The investment flowing into AI for ESG is not speculative. It's accelerating, and the numbers are steep.

2.1 AI in ESG Market Size[3]

The AI in ESG market is projected to grow from \$1.24 billion in 2024 to \$14.87 billion by 2034, a 28.2% CAGR. The generative AI segment dominates with 41.8% market share in 2024, and Data Collection & Analysis accounts for 37.3%.

2.2 Enterprise Investment Trends (KPMG Survey)[8]

KPMG surveyed 350 executives across 15+ countries on AI and sustainability. The signal is clear: 76% expect to increase Green IT investments in the next 2-3 years, 74% anticipate improved ROI, and 68% have already established strategic Green IT goals.[8] Leadership is funding this and expecting returns.

3. The Meta-Pattern: What the Big Four Actually Built

These firms have moved past proofs-of-concept into production systems that run core business processes at scale. They are not AI tools that answer questions when asked. They are agentic labor systems with governance: running processes end-to-end, making decisions within defined boundaries, escalating when they can't, and producing auditable output.

The underlying architecture shares five common characteristics:

1. **Role-based agents, not generic assistants** — Each agent has a specific domain and task scope
2. **Human oversight mapped to risk level** — High-risk decisions require human approval; low-risk actions execute autonomously
3. **Composable platforms** — Agents are orchestrated together, not invoked in isolation
4. **Multi-model interoperability** — Platforms support multiple AI providers (cloud-agnostic deployment)
5. **Built-in governance** — Compliance controls and review tracks embedded into the overall IT process

For sustainability teams, the infrastructure exists. The question is applying it to ESG-specific workflows.

4. Six Categories of AI Application — Mapped to Sustainability

The Big Four's deployments cluster into six categories. Each maps directly to a sustainability use case.

#	Category	What It Does	Sustainability Use Case
1	Deterministic Compliance	Rules-based checks, automated validation	Regulatory reporting, threshold monitoring
2	Risk Detection	Pattern recognition, anomaly flagging	Supply chain ESG risk, emissions anomalies
3	Knowledge Compression	Summarize, extract, synthesize documents	Supplier sustainability reports, policy documents
4	Analytical Synthesis	Cross-source analysis, trend identification	Emissions trends, benchmarking against peers
5	Workflow Orchestration	Multi-step automated processes	Supplier data collection, report generation
6	Product-Embedded Services	AI features in products/services	Dashboards, carbon accounting tools, ESG scoring

4.1 Deterministic Compliance Automation

Rules-heavy, high-repeatability tasks with automated validation. For ESG teams this means framework mapping and disclosure checklist completion across GRI, SASB, TCFD, CSRD, and CDP: the lowest-risk starting point, where rules are well-defined and the margin for error is measurable.

ESG Reporting Frameworks at a Glance

Framework	What It Is
GRI	Global Reporting Initiative. Most widely used global sustainability reporting standards.
SASB	Industry-specific, financially material disclosures.
TCFD	Climate risk and opportunity framework for investors.

Framework	What It Is
CSRD	EU mandatory reporting, affecting thousands of companies globally.
CDP	World's largest environmental disclosure system (23,000+ companies).

4.2 Risk Detection & Exception Intelligence

Probabilistic analysis with human-in-the-loop review. Turns sustainability from reactive ("we discovered a problem in last year's report") to proactive ("we caught the anomaly in real time"). Deloitte Zora AI includes real-time supplier ESG compliance evaluation as a core feature.[6]

4.3 Knowledge Compression & Expert Augmentation

Retrieval and reasoning over large knowledge bases. Sustainability teams are often small, sometimes one person, managing an expanding universe of frameworks and regulations. Knowledge compression lets a small team operate with the institutional knowledge of a large one.

4.4 Analytical Synthesis & Advisory Acceleration

Narrative and quantitative analysis across multiple sources: sustainability report drafting from structured data, materiality assessment synthesis, scenario modeling for decarbonization pathways. This is where AI moves from automation into insight generation.

4.5 Workflow Orchestration & Delivery Automation

Multi-step automated processes coordinating across departments and systems. ESG data is fragmented by nature: energy from facilities, procurement from finance, fleet from operations. Workflow orchestration solves the "ask 5 people" problem.

4.6 Product-Embedded AI Services

Unlike categories 1-5, which optimize internal processes, product-embedded services are the AI capability itself. Continuous emissions monitoring dashboards, carbon accounting tools, ESG scoring platforms. For climate tech startups, this is where the product opportunity lives: outcome-based pricing, not seat licenses.

5. AI Maturity Model for Sustainability Practitioners

Existing AI maturity models (Gartner, McKinsey, PwC, MITRE) are built for IT leaders and enterprise transformation teams. None address sustainability practitioners directly. None include sustainability as a dimension.

To fill that gap, I surveyed three established frameworks (McKinsey AI Readiness Index, PwC AI Maturity Assessment, MITRE AI Maturity Model), identified six common dimension clusters, reduced them to three core dimensions, and added Sustainability as a fourth — absent from all existing models.

Because dimensions 1-3 are domain-independent, the model supports cross-functional conversations about readiness. Only dimension 4 is sustainability-specific.

5.1 The Four Dimensions

Dimension	Core Question
1. Strategy & Leadership	Do you have a plan? Is the company aligned to it?
2. Data & Technology	Do you have the data infrastructure to support AI?
3. Organization & People	Does your workforce have the right training and governance?
4. Sustainability	Is the AI initiative enabling sustainability goals?

5.2 The Five Maturity Levels

Level	Name	What It Looks Like
1	Beginning	No plan, no skills, no tools. "We should look into that someday."
2	Experimenting	Individual initiative. Someone trying AI on their own. No formal support.
3	Operationalizing	Formal strategy exists. Budget allocated. Some workflows automated. Governance emerging.
4	Scaling	AI embedded in operations. Cross-functional. Multiple use cases in production. Audit-ready.
5	Optimizing	AI is assumed. Continuous improvement. Leading practices. Competitive advantage.

Not every organization needs Level 5. Target maturity depends on your mission, resources, and risk tolerance.

5.3 Dimension 1: Strategy & Leadership

Do you have a plan? Is the company aligned to it?

Level	What It Looks Like
1	"We should use AI for our ESG report... eventually." No timeline, no owner, no budget.
2	Sustainability manager experiments with ChatGPT for draft reports. Leadership unaware.
3	"AI for sustainability" is a line item in the annual plan. Budget allocated. VP sponsors a pilot.
4	CEO mentions AI + sustainability in earnings calls. CFO tracks ROI. Multiple funded projects.
5	Board asks "what's our AI sustainability strategy?" AI embedded in all sustainability planning cycles.

5.4 Dimension 2: Data & Technology

Do you have the data infrastructure to support AI?

Level	What It Looks Like
1	Emissions data in 12 spreadsheets across 5 departments. "Ask finance for electricity, ask ops for fleet."
2	Using ChatGPT to summarize supplier sustainability reports. Copy-paste from PDFs.
3	Sustainability data warehouse exists. Approved access to AI tools. Automated collection from some sources.
4	ESG report pulls from live pipelines. AI analyzes trends across all facilities.
5	Ask any sustainability question, get an answer. Predictive models for emissions. Real-time dashboards.

5.5 Dimension 3: Organization & People

Can your people use AI responsibly?

Level	What It Looks Like
1	"AI? That's IT's thing." No AI skills on the sustainability team. No governance policy.

Level	What It Looks Like
2	One analyst self-taught on ChatGPT. Shares tips informally. No formal review process.
3	Team attended AI workshop. Most comfortable with AI tools. Governance policy covers sustainability use. Human reviews AI content before publication.
4	AI skills in job descriptions. Sustainability and data teams meet weekly. Can explain AI methodology to auditors.
5	Sustainability team trains other departments on AI for ESG. Audit-ready AI governance.

5.6 Dimension 4: Sustainability

Is the AI initiative enabling sustainability goals?

This is what makes the model specific to sustainability. None of the existing frameworks assess whether AI is actually improving sustainability outcomes.

Level	What It Looks Like
1	No AI in sustainability work. Manual data collection, Excel reports.
2	Using ChatGPT to draft report sections or summarize supplier documents. Ad-hoc, individual initiative.
3	AI integrated into sustainability workflows. Automating data collection from some sources. Starting to generate insights beyond reporting.
4	AI across sustainability operations. Automated carbon accounting. Predictive models for emissions. Tracking own AI carbon footprint.
5	AI is how sustainability work gets done. Real-time dashboards. Ask any sustainability question, get an answer. Leading industry practices.

5.7 The Gap

Most ESG teams sit at Level 1-2 across the board, while finance and audit at large enterprises are at Level 3-4. The same AI platforms and patterns that got finance there can be applied to sustainability work. You don't have to build from scratch. You need to get in the room where these tools are being deployed and make sure sustainability is on the agenda.

6. Implementation Roadmap

6.1 For Enterprises: Start with Data, Then Intelligence

Most enterprises share the same fundamental problem: sustainability data is scattered across departments, systems, and spreadsheets. The sequence below reflects that reality. Each phase builds on the one before it.

Phase 1: Data Infrastructure Solve data fragmentation first. Agents collect, normalize, and reconcile across departments. Without clean, reconciled data, nothing else in this sequence works.

Phase 2: Supplier & Scope 3 Intelligence The biggest blind spot for most organizations. Bridge procurement and sustainability data. This is where the highest strategic value lives: Scope 3 typically represents 70-90% of an organization's total carbon footprint.

Phase 3: Reporting & Assurance Follows naturally once data is clean. Automate framework mapping across GRI, SASB, TCFD, and CSRD. Prepare workpapers for audit and assurance.

Phase 4: Continuous Management Move from annual scramble to continuous system. Real-time monitoring, anomaly alerts, and live dashboards. This is the long-term operating model shift.

6.2 For Climate Tech Startups: Build Products, Not Internal Tools

Climate tech startups and corporate innovation groups can build AI-native from day one. The opportunity is to productize what enterprises struggle to build internally. CSRD is driving immediate demand for ESG Reporting as a Service — framework mapping, disclosure drafting, assurance prep. The Scope 3 and supplier intelligence gap is even larger: procurement data to emissions estimates, with supplier monitoring built in. And ESG Agent as a Service, with outcome-based pricing rather than seat licenses, is how the Big Four are already shifting their own business model. Startups can build this way from day one.

6.3 Risk Callout: Not All AI Projects Succeed

Before planning your implementation, the headwinds deserve an honest look. Gartner projects 40% of agentic AI projects will be canceled by end of 2027 due to escalating costs, unclear business value, or inadequate risk controls.[9] Only around 130 of thousands of "agentic AI vendors" are considered legitimate — the rest is agent washing.[9] And less than 5% of enterprise apps have AI agents today, though Gartner projects that rising to 40% by 2026.[10]

The Big Four have resources to absorb failures. Most organizations don't. Start with proven patterns and the highest-pain workflow you can identify.

7. The Scope 3 Problem: Why Procurement is the Choke Point

Scope 3 emissions typically represent 70-90% of an organization's total carbon footprint, and they're the hardest to measure. The challenge is practical:

1. Data is fragmented across procurement, finance, and supplier systems
2. Suppliers don't report consistently, or at all
3. Estimation methods vary and are hard to verify
4. Procurement teams own the relationships but not the sustainability expertise

The regulatory picture in the US is mixed: the SEC dropped Scope 3 from its climate disclosure rule, but California's SB 253 requires Scope 3 reporting starting in 2027 for companies with \$1B+ revenue doing business in the state. The EU's CSRD is phasing in Scope 3 requirements for companies with European operations. Even where regulation is uncertain, market pressure is real. Large customers and institutional investors are increasingly requesting supply chain emissions data.

AI agents that bridge procurement data (spend, contracts, supplier info) with sustainability calculations (emissions factors, ESG scores) address this directly. Deloitte Zora AI already includes real-time supplier ESG compliance evaluation and spend analysis in one platform.[6]

For sustainability leaders: Your procurement team is your most important ESG partner. The agents that help them are the agents that help you.

8. AI's Environmental Footprint

Using AI to track environmental impact while that AI itself has an environmental footprint is a real tension, consistently raised by practitioners. A single ChatGPT query consumes roughly 0.34 Wh,[11] but complex reasoning queries can use up to 50x more, and inference — not training — dominates total AI energy at scale.[13] The carbon impact varies enormously by where workloads run: the same activity produced 20x more CO₂ on the US average grid than on France's nuclear-heavy grid.[14] Water consumption is the most underreported dimension, and it warrants explicit consideration for organizations with water reduction commitments.[12]

The IEA's independent analysis concludes that AI's emissions reduction potential in end-use sectors is 3-4x larger than its own data center footprint by 2035.[15] Production results back this

up: from 40% reductions in data center cooling energy[16] to materials discoveries accelerating clean energy technology.[17]

For a full treatment of measurement tools, vendor disclosure gaps, cloud provider dashboards, and the FinOps-to-GreenOps bridge, see the companion paper: *The Sustainable AI Paradox: What Sustainability Leaders Need to Know About AI's Environmental Footprint* (Wynne Technologies, March 2026).

9. Strategic Takeaway

ESG is following the same path as audit and tax. The Big Four have proven the architecture works. ESG reporting is still manual, spreadsheet-driven, and fragile under audit — but agents can turn that annual scramble into a continuous system. Scope 3 is the hardest problem and procurement is the choke point. The tools that solve it already exist. They're just running under different names in audit and tax today.

Where to start: Use the maturity model in Section 5 to assess where your organization stands across the four dimensions. Most sustainability teams are at Level 1-2. What would it take to get to Level 3? That question, and the highest-pain workflow you can identify, is your starting point.

10. Next Steps

Applying these patterns to your organization requires five things:

1. **Maturity assessment** — Where does your organization sit across the 4 dimensions?
 2. **Use case prioritization** — Which of the six categories fits your highest-pain workflows?
 3. **Procurement alignment** — Are sustainability and procurement working together on Scope 3?
 4. **Platform evaluation** — Build, buy, or partner?
 5. **Governance framework** — How will you maintain human oversight and manage AI's own footprint?
-

About the Author

Adam Wynne is the founder of Wynne Technologies, where he advises on company-wide AI technical strategy and helps product teams ship faster by combining product discipline with AI-accelerated development. After 20 years building SaaS products for organizations including

Bosch, Armada Supply Chain Solutions, and the US Department of Energy's Pacific Northwest National Laboratory (PNL), he now applies that expertise to climate challenges and sustainability technology.

Learn more at wynnetech.ai.

Footnotes

[1] PwC Global News Room, "PwC scales global AI agent ecosystem," October 2025.

<https://www.pwc.com/gx/en/news-room/press-releases/2025/scales-global-ai-agent-ecosystem.html>

[2] EY Value Realized 2025 Report. EY states "1,000 AI agents in development or production" with target of "100,000 AI agents by 2028."

[3] Market.us, "AI in ESG and Sustainability Market Size," 2024.

<https://market.us/report/ai-in-esg-and-sustainability-market/>

[4] EY Newsroom, "EY launching EY.ai Agentic Platform," March 2025.

https://www.ey.com/en_us/newsroom/2025/03/ey-launching-ey-ai-agentic-platform

[5] KPMG Press Release, "KPMG launches multi-agent AI platform," June 2025.

<https://kpmg.com/xx/en/media/press-releases/2025/06/kpmg-launches-a-multi-agent-ai-platform>

[6] Deloitte Press Room, "Deloitte unveils Zora AI," March 2025.

<https://www.deloitte.com/us/en/about/press-room/deloitte-unveils-zora-ai-agentic-ai-for-tomorrows-workforce.html>

[7] PwC Newsroom, "PwC launches AI Agent Operating System for Enterprises," March 2025.

<https://www.pwc.com/us/en/about-us/newsroom/press-releases/pwc-launches-ai-agent-operating-system-enterprises.html>

[8] KPMG, "AI and Sustainability Survey," 2025. <https://kpmg.com/us/en/articles/2025/ai-sustainability.html>

[9] Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," June 2025.

<https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>

[10] Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," August 2025.

<https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026>

[11] OpenAI, 2025. Via MIT Technology Review, "The carbon footprint of AI," May 2025. ChatGPT consumes approximately 0.34 Wh per query; reasoning models can use up to 50x more.

[12] Li, P., Yang, J., Islam, M.A., Ren, S. "Making AI Less Thirsty: Uncovering and Addressing the Secret Water Footprint of AI Models," Communications of the ACM, 2023. GPT-3 consumes 500ml of water per 10-50 medium-length responses depending on deployment location. <https://arxiv.org/abs/2304.03271>

[13] Deloitte, "Technology, Media & Telecommunications Predictions 2025." Projects AI inference to account for roughly 66% of AI energy consumption by 2026.

[14] Luccioni, A.S., Viguier, S., Ligozat, A. "Estimating the Carbon Footprint of BLOOM," 2023. BLOOM training on French nuclear grid: 25 tCO₂e. Patterson et al., "Carbon Emissions and Large Neural Network Training," 2021. GPT-3 training on US average grid: 502 tCO₂e. <https://arxiv.org/abs/2211.02001>

[15] IEA, "The energy impact of AI," 2025. <https://www.iea.org/topics/artificial-intelligence>

[16] DeepMind blog, "AI for Data Centre Efficiency," 2016/2018. <https://deepmind.google/discover/blog/>

[17] Merchant, A. et al., "Scaling deep learning for materials discovery," Nature, 2023. GNoME discovered 2.2 million new crystal structures relevant to batteries, solar cells, and carbon capture. <https://www.nature.com/articles/s41586-023-06735-9>