



Overview— The IT Toolkit for Responsible and Sustainable AI

A Field Guide for
Implementation at Scale

By Richard Pastore
Research Principal
SustainableIT.org

**IT Toolkit for
Responsible &
Sustainable AI:**
A Field Guide for
Implementation
at Scale



Advanced artificial intelligence technologies available to enterprises today offer revolutionary innovation opportunities. Generative AI systems and their Large Language Models are transforming the creation of content with unprecedented speed and scope, while Agentic AI embeds autonomous decision-making and execution into revolutionized workflows.

However, the design and use of these tools pose considerable risks for the environment, equity, trust, security, and privacy. To optimize the benefits and minimize the risk of negative impacts, organizations must operationalize the principles and practices to govern AI systems responsibly and sustainably throughout their lifecycles.

The term “Responsible AI” refers to the governance and operational discipline that ensures AI systems are developed, deployed, and maintained in ways that align with legal, ethical, and organizational standards. It includes oversight practices, risk management, compliance, and lifecycle accountability across systems and teams. “Sustainable AI” is differentiated by its focus on environmental impact, including the limiting of non-renewable energy consumption, water usage, and carbon emissions footprint. It also incorporates social and governance sustainability.

Guiding and operationalizing responsible transformation isn’t simple, or even common in the case of AI. Responsible governance is assumed to be an impediment to AI uptake. Yet, IT leaders with responsible AI platforms say they accelerate deployment of effective AI models and reduce time wasted on dead ends. The IT Toolkit for Responsible and Sustainable AI was developed by a SustainableIT.org working group of global technology, business, and AI experts to provide executives with a strategic framework and practical guidance for navigating this transformation. Grounded in real-world practices and aligned with emerging standards, it equips organizations to not only mitigate AI-related risks but also unlock lasting business value and innovation at scale.

The Toolkit offers the following inaugural resources. These will be augmented and updated as warranted.

1. **Reference Guide to Responsible AI Governance Frameworks, Principles, and Standards**
A curated reference of leading public resources and standards that inform responsible AI governance strategies
2. **AI General Glossary** Explanations of terms and buzzwords
3. **Responsible AI Leadership Framework and Principles** A foundational framework of nine key principles for responsible and sustainable AI lifecycle leadership
4. **Responsible AI Lifecycle Model** A model with mapping responsible governance across the combined lifecycles of AI development and DevOps
5. **AI Sustainability Runbook – A Practical Guide**
Hands-on guidance to operationalize sustainability within AI systems, ensuring they support long-term ESG goals throughout their lifecycle
6. **Responsible AI Data Governance Principles & Runbook – A Practical Guide**
Step-by-step guidance for managing AI data with integrity, transparency, and environmental responsibility—based on five focusing principles for sustainable AI data practices
7. **Persona-based Literacy Guide for Sustainable AI**
Tailored learning paths to help key enterprise personas build the knowledge needed to develop, govern, and apply AI sustainably
8. **Best Practices and Strategic Takeaways from the Responsible AI Impact Awards**
Summaries of award-winning AI-based applications, platforms, products, and services that generate high-impact business value while exemplifying responsible and sustainable lifecycle best practices
9. **Generative AI Ecosystem Model for IT**
The model emphasizes that AI's environmental impact extends far beyond the moment of use (inference), encompassing manufacturing, training, infrastructure, and end-of-life.

The Toolkit is designed as a practical, enterprise-wide resource to help organizations govern and operationalize AI responsibly and sustainably across the full lifecycle—from ideation to impact. Start with the fundamentals—the Leadership Framework and Principles, Governance Lifecycle Model, and AI Literacy Primer all outline key concepts for enterprise leaders and their teams. The Glossary and Reference Guide are there for background knowledge.

The more granular Sustainability and Data Governance Runbooks delve into critical steps to operationalize AI responsibility and sustainability, with supporting KPIs, documentation, templates, and links. For inspiration and real-world implementation examples, explore the Responsible AI Impact Award winner descriptions.

The Case for Responsible AI

Most AI system users have little to no idea of the energy consumption and emissions instigated by their prompts, nor do they anticipate the biases inherent in their data and propagated by their outputs. AI's risks and challenges—ethical, social, and most particularly environmental—have been left out of much of the hype for these tools. The following is a brief rundown of the dimensions of some of AI's primary business risks.

- » **Ethics** – The use of GenAI's creative abilities raises rights and ownership dilemmas for intellectual property and data. The misuse of GenAI, for deep fakes and other forms of misinformation, can cause personal and societal harm.
- » **Digital Divide** – Without broad inclusivity and equitable access to AI applications and their benefits, the technology could lead to a widening of society's digital divide. Approximately 2.6 billion people globally remain without internet access, a fundamental barrier to equitable access to AI technologies.¹
- » **Integrity** – AI decisions are susceptible to inaccuracies and discriminatory outcomes due to biases in data. There is a risk of generating low-quality or biased output, which can harm business reputations and erode customer and institutional trust.
- » **Resource Intensity** – AI data centers are significant energy consumers. For instance, data centers in Phoenix consume approximately 1.5 gigawatts of power.² This is approximately the same amount of energy that would power 1.2 million U.S. homes for a single day, or a city the size of Scottsdale, Arizona for over 10 days. The cooling requirements of AI data centers lead to high water consumption. ChatGPT, for example, uses about two liters of water for every 10 to 50 queries. At the high end of the range, this would be enough water to fill 16 Olympic-sized swimming pools every day.³
- » **Climate Degradation** – The emissions generated in the GenAI lifecycle, from training through usage, are staggeringly high compared to other technologies. Since 2012, the most extensive AI training runs have been using exponentially more computing power, doubling every 3.4 months, on average.⁴
OpenAI's GPT-3 training is estimated to have used 1.3 gigawatt-hours of energy (equivalent to 120 average U.S. households' yearly consumption) and generated 552 tons in carbon emissions (equivalent to the yearly emissions of 120 U.S. cars).⁵ OpenAI's GPT-4, is said to be 10 times larger.⁶

Job Displacement – While AI can augment human roles, there is also concern about large-scale job displacement, particularly in industries where routine tasks are easily automated. This requires businesses to consider retraining and upskilling initiatives to avoid social harm.

Security and Privacy Risks – AI can be exploited for malicious purposes, such as creating fake identities or generating harmful content. Use of data to train and inform AI algorithms may inadvertently violate privacy laws or regulations. Businesses must invest in the latest security measures to protect against these risks.

If avoidance of these negative impacts does not inspire urgency to apply a responsible approach to AI, then the upside of a sustainable AI portfolio should be an added incentive. By stepping forward to propose and implement the responsible strategies and practices contained in the Toolkit, IT organizations will directly contribute to the future competitiveness of their enterprises. Essential benefits include:

¹ Reuters, Comment, *Together we can end the digital divide that disenfranchises 2.6 billion people*, September 2024

² Stanford University, *Thirsty for power and water, AI-crunching data centers sprout across the West*, <https://andthewest.stanford.edu>

³ The Times, *'Thirsty' ChatGPT Uses Four Times More Water Than Previously Thought*, October 2024

⁴ AI and Compute, OpenAI, May 16, 2018

⁵ Patterson etc., *Carbon Emissions and Large Neural Network Training*, April 2021

⁶ M. Schreiner, *The Decoder, GPT-4 Architecture, Datasets, Costs and More Leaked*, July 2023

1 0 1

1 1

- » **Increased Confidence in AI Technology** – Implementing ethical usage principles aligns AI applications with the company's values and ethical standards, increasing corporate confidence in AI-derived decisions and actions.
- » **Customer Trust and Loyalty** – AI outcomes that are fair, transparent, and explainable help build corporate, customer, and institutional trust in business decisions. This will help maintain long-term customer loyalty and company reputation.
- » **Advancement of Environmental Goals** – Conducting environmental sustainability due diligence not only advances the company's sustainability goals, but it improves competitiveness in markets where environmental stewardship is—or will soon become—highly valued.
- » **Future-proofed Workforce** – By prioritizing AI literacy and human-first AI principles, companies equip their current and future workforce to engage effectively with AI technologies.
- » **Optimized Innovation and Efficiency** – By adhering to a responsible AI governance framework, companies can more freely tap AI's potential without the distrust, delays, fines, and other consequences of poorly governed technology rollouts.

Responsible AI Needs a Solid Foundation

These responsible AI benefits will accrue only if organizations support AI governance with fundamental strategies and practices. Savvy leaders will always start by asking a deeper question than “Which AI should we deploy?” They first consider whether the workflow deserves automation at all, then redesign the work so the chosen AI technology will augment human strengths rather than embalming an obsolete process. As Ben Blanquera, vice president of technology and sustainability at Rackspace Technology, writes:


Most AI initiatives don't fail because the technology doesn't work. They fail because the process didn't deserve automation in the first place.

The industry is obsessing over how to implement AI—RAG pipelines, LLM fine-tuning, vector databases. But the real question is: WHAT are we trying to automate or augment? The top performers in AI aren't just tech-forward—they're imagination-forward. They rethink the work itself. They analyze the value chain with both creativity and rigor.

Before deploying GenAI, copilots, agents, or automation, we need to step back:

- » What tasks should disappear entirely?
- » What flows should be restructured from scratch?
- » Where do human strengths still offer unique value?

AI isn't the strategy. Reimagining work is.



At the same time that leaders are evaluating whether AI is right for them, they should ask if the organization is ready for AI. As Thomas Heine, publisher of the magazine for sustainable IT procurement “Kleine Kniffe” and CEO of SDG Media GmbH, writes:

You’re in a strategy meeting. A new AI tool is on the agenda—powerful, efficient, scalable. The promise is clear: faster decisions, better predictions, smarter processes. But alongside the potential, one key question arises: Is our organization truly ready for this?

AI is no longer a topic for the future. It’s reshaping business models, markets, and day-to-day operations across all levels. Today’s leaders need more than just access to new tools—they need a clear understanding of how AI works, where its strengths lie, and where it must be questioned.

Organizations need a foundation for sound, strategic judgment: What can AI really do? Where are its limits? What are the critical questions we must ask before integrating it into our processes, products, or strategies?

AI literacy doesn’t mean knowing how to code. It means being able to lead informed conversations. To spot risks. To identify opportunities. And above all, to guide teams, provide clarity, and take responsibility for using AI in a meaningful, accountable way.

Organizations that invest in AI literacy are investing in resilience and future-readiness. Because those who understand AI make better decisions—today and tomorrow.

Another fundamental enabler for a sustainable AI portfolio is to integrate AI-specific sustainability and ethics checkpoints into IT’s operating models, like DevOps. Sustainability principles—such as minimizing carbon intensity, using energy-efficient models, and monitoring model drift—become part of the automated testing and deployment stages. For example, pre-deployment gates can assess whether AI models meet predefined thresholds for compute efficiency (e.g., gCO₂e/inference), data integrity, and explainability. DevSecOps can be extended to include “DevSustAI” practices, ensuring that AI systems are not only secure and compliant but also environmentally and socially responsible throughout the development lifecycle.

Similarly, FinOps, which emphasizes financial accountability and optimization in cloud environments, can evolve to include environmental impact metrics alongside cost. Aligning compute usage with sustainability goals could include selecting cloud regions powered by renewable energy, evaluating the ROI of smaller, more efficient AI models, and factoring in sustainability KPIs like emissions per dollar spent or energy per model inference. Embedding sustainability into FinOps bridges financial and ESG objectives, enabling enterprises to make AI deployment decisions that are cost-effective, efficient, and aligned with long-term responsibility commitments.

Responsible AI must also become embedded in enterprise architecture—the business and tech blueprint for ongoing operations. As Lisa Pratico, chair of Sustainable Architectures and Responsible AI for IASA Global, writes:

Enterprise architecture (EA) plays a critical role in ensuring that Responsible AI initiatives are aligned with corporate sustainability objectives, ESG mandates, and long-term innovation goals. Through structured strategy enablement and business-case rigor, EA helps organizations prioritize ethical, energy-efficient, and scalable AI adoption.

EA's contributions to responsible AI strategy are numerous:

- » **Strategic Alignment** – EA ensures AI use cases and technology choices align with environmental goals (e.g., carbon reduction), social equity (e.g., bias-free AI), and governance standards (e.g., auditability, transparency).
- » **Architecture Principles** – EA codifies sustainability-focused design standards by leveraging industry frameworks such as TOGAF, IASA Global, SDIA, NIST AI RMF, and ISO/IEC 42001.
- » **Scenario Evaluation** – EA models the environmental and ethical implications of AI options, using capability maps and solution blueprints to support or redirect AI investments.
- » **Governance Integration** – EA embeds sustainability checkpoints into Architecture Review Board (ARB) deliberation, ensuring all new AI initiatives are evaluated for energy impact, carbon lifecycle footprint, data ethics and fairness, and resource efficiency.

In certain situations, AI model selection may be vetoed or redirected based on projected emissions or bias impact assessments as modeled through EA frameworks. EA also helps ensure that business cases for AI-enabled systems embed responsibility and sustainability.

With a solid architectural, literacy, and operational foundation in place, organizations will be ready to develop, adapt and scale transformative AI solutions. Leveraging the practical guidance of The IT Toolkit for Responsible and Sustainable AI, IT leaders and their business peers will harness AI in ways that are not only powerful and profitable, but also principled and sustainable. Now is the time to act. Dive in, share it across your teams, and embed these practices into your AI programs. A sustainable future for AI is not just possible, it's within reach.



Contributors

We at SustainableIT.org are indebted to the volunteers that devoted time from January through June, 2025, to scope, develop, critique, and iterate the inaugural Toolkit resources. Their personal commitment to the cause has been inspirational and they have been a pleasure to work with.

Responsible AI Working Group Leaders

- » **Wiebren van der Zee**, Sustainability Domain Leader, Office of the CIO, ABN AMRO Bank
- » **Berndt Pilgram**, Lead Product Data & Information Security Governance, Fresenius Medical Care
- » **Ben Blanquera**, VP, Technology and Sustainability, Rackspace Technology
- » **Jody Elliott**, Head of IT Risk and Sustainability, National Grid
- » **Ron Guerrier**, CTO, Save the Children
- » **Lisa Pratico**, Chair, Sustainable Architectures and Responsible AI, IASAGlobal.org
- » **Christopher Caldwell**, Senior Business Leader DevOps, Mastercard
- » **Justin Mann**, Head of Research & Standards, TBM Council

Project Leader

- » Richard R. Pastore, Research Principal, SustainableIT.org

Workstream Contributors

- » **Shane Artis**, Distinguished Engineer, Head of Sustainable Finance and Resilience Technology, Morgan Stanley
- » **Jack Bischof**, General Manager and EME Region VP, TBM Council
- » **William E. Buchman**, Senior IT Officer, Information Management Services, World Bank Group
- » **Katia Chaban**, VP, Global Head of IT Sustainability Center of Expertise, NTT Data
- » **Ben Christensen**, Co-Chair, Climate, AI2030
- » **Kacy Clarke**, Co-founder, Sustainable Architectures and Responsible AI, IASAGlobal.org
- » **Marie Daille**, Responsible Technology Leader, Amer Sports
- » **Angelo Fienga**, Senior Business Advisor, Deloitte

- » **Stacey Fleckles**, Head of Audit, Risk and Compliance—Digital Platforms and Infrastructure Technology, National Grid
- » **Anshuman Ghosh**, Principal Process Consultant, Tech Strategy & Innovation, Danske Bank
- » **Jeanne Goncalves**, Senior IT Officer, Client Services, World Bank Group
- » **Saranthan Govindarajan**, Former VP and CTO, Mattel
- » **Sunil Gupta**, Head, Data Science Center of Excellence, Blue Cross Blue Shield North Carolina
- » **Hannah Homstedt**, Responsible GenAI Expert, Danske Bank
- » **Valence Howden**, Advisory Fellow, Info-tech Research Group
- » **Raphael Jung**, MSc Strategic Product Design, Delft University of Technology
- » **Rainer Karcher**, Global Head of Sustainability, Heartprint GmbH, and CSO, SustainableIT.org
- » **Ivneet Kaur**, Former Executive VP and Chief Information and Technology Officer, Sterling
- » **Despoina Kouria**, Sustainability Advisor, World Bank Group
- » **Srini Koushik**, President, AI and Sustainability, Rackspace Technology
- » **Borislava Koycheva**, Head of Sustainable IT and Strategy Principal, Volvo Group
- » **Jerry Lassoie**, Vice President, Data Center Engineering EMEA, Mastercard
- » **Sébastien Le Callonnec**, Sr. Principal Software Engineer, Mastercard
- » **Ralph Loura**, Board Chairman, SustainableIT.org
- » **Sharon Mandell**, CIO, Juniper Networks
- » **Wendy Marshall**, Regional Technology Lead, Rackspace Technology



- » **David Ray**, Senior IT Risk Management Officer, World Bank Group
- » **Sherri Reardon**, Director of Client Relations, SustainableIT.org
- » **Anurag Revri**, Chief Enterprise Architect, Stanford Health Care
- » **Mayya Revzina**, Senior Publishing Officer, World Bank Group
- » **Mark Richter**, Managing Principal, Calais Cloud
- » **Anne Rosenberg**, Former Executive VP, Strategy and Global Expansion, SustainableIT.org
- » **Fabian Schnippering**, Sustainable IT Team Lead, Capgemini Invent
- » **Oscar Slotosch**, CEO, Validas AG
- » **Lasse Borris Sørensen**, Founder, Plan A Projects
- » **Melissa Stires**, Head of Global Growth & Expansion, Mia AI
- » **Chitra Sundaram**, Data and AI Track Lead, SustainableArchitectures.org
- » **Niklas Sundberg**, Chief Digital Officer and SVP, Kuehne + Nagel
- » **Craig Vezina**, CEO, ThespaceShip.org
- » **Sruthi Viswanathan**, CTO and Co-founder, The Spaceship Academy
- » **Jens Wandel**, Special Advisor to the UN Secretary-General on Reforms, United Nations
- » **Bill Wong**, AI Research Fellow, Info-Tech Research Group
- » **Ann-Marie Yap**, Executive Director of Technology, Stanford Health Care
- » **Marisa Zalabak**, CEO and Founder Open Channel Culture; Co-chair, IEEE AI Ethics Education & Chair, Planet Positive 2030
- » **Xiaochen Zhang**, Chief AI Officer & Executive Director, AI2030

Designed by Jennifer Wanderer, Art Director, Jenniferwanderer.com

About SustainableIT.org

Vision

Advancing global sustainability through technology leadership.

Mission

Our mission is to unite the world's largest community of technology and sustainability leaders to define sustainability transformation programs, author best practices and frameworks, set standards and certifications for governance, provide education and training, and raise awareness for IT-centric ESG programs that make their organizations and the world sustainable for generations to come.

Mandates

Best Practices, Research and Standards

Identify sustainable digital transformation programs by industry. Research and define best practices, frameworks, and standards for all three pillars of sustainability (environmental, societal, and governance) for IT departments and organizations.

Global Awareness and Recognition

Promote sustainable digital transformation programs and advances in sustainability. Raise awareness through local, regional, and global awards, as well as through social media, publications, and public relations.

Community, Education and Training

Build local and regional communities for technology leaders to advance sustainability. Develop education and training programs for IT leadership and professionals for all three pillars of sustainability.

Transparency and Accountability

Set standards for metrics and reporting to enable transparency and accountability. Create certification programs for individuals and organizations with rights to use our sustainability emblem.



About SustainableIT.org

SustainableIT.org is a Delaware 501(c)(6) nonprofit, non-stock legal entity led by technology executives who will advance global sustainability through technology leadership. Our mission is to define sustainable transformation programs, author best practices and frameworks, set standards and certifications, provide education and training, and raise awareness for environmental and societal programs that make our organizations and the world sustainable for generations to come.