

10 Actions for Using AI in Business

Grounded in Regenerative AI Ethics - a framework aligned with Life's principles

The AI train has left the station. But **who is driving**, and where is it heading?

This guide gives business leaders 10 concrete, practical actions grounded in Regenerative AI Ethics, a framework built not just on human safety, but on the principles that have sustained Life on Earth for 3.8 billion years.

You do not need to be a computer scientist to use this guide. You need to be someone who cares about your people (customers, employees, suppliers, etc.), your community, and the world your children will inherit.

Martina Doleshal

Co-Author, Regenerative AI Ethics Framework (<https://www.regenerativeaiethics.com>)

Systems Thinker, Futurist and Serial Entrepreneur (<https://doleshalconsulting.com>)

How to use this guide

Each of the 10 actions that follow is structured in four parts:

- **Why this matters** - the systemic reason this action is necessary, grounded in research and regenerative principles.
- **The core action** - what to do, how to implement the initiative.
- **What this looks like in practice** - illustrative examples from real organisational contexts.
- **Actions to take this month** - specific next steps you can begin immediately.

This is not a theoretical document. Every action here can be started without a large budget, without waiting for regulation, and without a dedicated AI team. It requires intentionality, leadership, and the willingness to ask harder questions than your industry currently expects.

A note on framing. Conventional AI ethics focuses on avoiding harm. Regenerative AI Ethics goes further as it asks how AI can actively contribute to the flourishing of people and planet, not just avoid damaging them. That is the lens these 10 actions are written through.

The 10 actions are grouped roughly from **foundational** (internal clarity and governance) through **operational** (vendor relationships, metrics, assessment) to **systemic** (culture, regenerative use, collective advocacy). Work through them in order, or start where your organisation has the most urgent need.

1

Start with a Values Audit Before Any AI Decision

WHY THIS MATTERS FIRST

Most organisations deploy AI tools to solve operational problems and only discover the ethical ones later. A values audit reverses that sequence, turning your existing commitments into a procurement filter before any contract is signed or any AI tool is switched on.

Before deploying any new AI tool, ask your team:

- What do we actually care about as a business and as people?
- What does 'good' look like for our employees, our customers, our community, and the planet?

Use these values as a filter to evaluate any AI tools you are considering for implementation in your business.

WHAT THIS LOOKS LIKE IN PRACTICE

A recruitment firm runs a 90-minute all-hands session asking: 'What kind of employer do we want to be known as in 10 years?' The answers became the explicit criteria against which every AI-assisted hiring tool was evaluated before demo calls were booked.

A retail business creates a one-page 'AI values card' with five sentences describing how AI must treat customer data, frontline workers, and suppliers. New tools are assessed against it at the shortlisting stage, not after purchase.

A healthcare provider adds 'does this align with our duty of care?' as a standing question on their technology procurement checklist, prompted by a concern that efficiency-optimizing scheduling tools were quietly deprioritising complex patients.

"If an AI tool optimizes for something that contradicts your values, it is not a bargain, it is a liability."

ACTIONS TO TAKE THIS MONTH

1. Host a 90-minute values workshop with a cross-section of your team, not just leadership.
2. Translate the output into 5–7 plain-language statements about what 'good AI use' means in your context.
3. Make those statements a required section of any AI procurement brief or vendor evaluation scorecard.
4. Revisit the audit annually or whenever your strategic direction shifts.

2

ACTION 2 OF 10

Conduct a Full AI Inventory

WHY YOU PROBABLY NEED THIS MORE THAN YOU THINK

Most leadership teams significantly underestimate how many AI-powered tools are already active in their organisation. AI has been embedded quietly into HR platforms, customer service software, accounting tools, and content systems for years often without anyone labeling it as AI or establishing oversight.

Map every AI tool currently in use across your organization.
For each one, document:
**what problem it solves, what data it uses,
what decisions it influences, and what oversight exists.**

WHAT THIS LOOKS LIKE IN PRACTICE

A 40-person marketing agency discovered 14 distinct AI tools in use after conducting their first inventory. Including tools making automated decisions about client ad spend, content moderation on client social accounts, and freelancer performance scoring. Only three had any documented oversight.

A financial services firm used their inventory to map which AI systems were influencing customer-facing decisions. They identified two tools producing outputs that differed systematically by postcode, a red flag for algorithmic discrimination that had gone unnoticed for 18 months.

A manufacturing company discovered their HR platform had enabled an 'AI scoring' feature by default in a software update. No one had knowingly activated it, but it had been ranking job applicants for three hiring cycles.

"Most businesses discover they are using many more AI tools than their leadership realizes and with far less oversight than they thought."

ACTIONS TO TAKE THIS MONTH

1. Survey all department heads: list every software tool that uses AI, automation, or algorithmic decision-making.
 2. For each tool, document: vendor, purpose, data inputs, outputs, decision influence, and who reviews results.
 3. Flag any tool influencing hiring, promotion, customer pricing, credit, or content as these require closer scrutiny.
 4. Establish a standing update process: any new AI tool must be added to the inventory before go-live.
-

3

Ask the Hidden Cost Questions

THE INVOICE YOU HAVEN'T SEEN

The price of AI is not just your SaaS subscription. Training large AI models consumes enormous energy. For example, running a single GPT-class model can use as much electricity as 130 U.S. homes for a year. By 2027, global AI water use for cooling data centers could equal six times Denmark's annual water consumption (OECD, 2023).

And only **12% of executives currently measure any of the AI impacts** (Capgemini, 2025).

Request environmental and labor impact information from your AI vendors. Ask about energy consumption, water use, carbon footprint, and the working conditions of data labelers. Compare models where possible as some platforms offer more energy-efficient options. Begin tracking your organisation's AI-related environmental footprint as a regular metric.

WHAT THIS LOOKS LIKE IN PRACTICE

A B-Corp certified food business added three questions to their standard vendor due diligence form: (1) What is the estimated energy footprint of this model? (2) What percentage of your training data workforce is directly employed vs. contracted? (3) Do you publish a sustainability report for AI operations? Most vendors couldn't answer, which became a selection criterion.

A mid-size law firm switched from one AI legal research platform to a competitor after discovering the first had no renewable energy commitments for its data centers. The switch required a 3-week transition but aligned with their ESG reporting obligations.

A consultancy began including an 'AI energy line' in their quarterly carbon reporting estimated from usage volume and vendor-published efficiency data. It created internal accountability that had simply not existed before.

"Being among the 12% of companies that measure AI's environmental footprint is both a competitive and a reputational advantage."

ACTIONS TO TAKE THIS MONTH

1. Add environmental and labour questions to every AI vendor RFP and renewal conversation.
2. Ask vendors for their data centre energy source mix and any renewable energy commitments.
3. Request information about the working conditions of human reviewers who label or moderate training data.
4. Start tracking AI-related energy use (even a rough estimate) in your sustainability reporting.
5. Explore whether lighter, more efficient model options (smaller models, cached queries) can meet your needs.

4

Build an AI Ethics Working Group

DECISIONS MADE IN SILOS CREATE BLIND SPOTS

AI decisions made solely by IT or by senior leadership routinely miss the concerns of the people closest to the impact, such as customer-facing staff, HR teams, and those from communities historically underserved by automated systems.

Diverse, cross-functional governance is not a bureaucratic overhead, but rather it is how organisations catch risks before they become crises.

Designate a cross-functional team, including people from HR, operations, legal, sustainability, and frontline roles, to guide AI adoption decisions.

This group should meet regularly, review new tool deployments, and be empowered to say 'not yet' or 'not this way.'

WHAT THIS LOOKS LIKE IN PRACTICE

A retail chain formed a five-person AI Ethics Panel including their head of sustainability, a frontline store supervisor, their data protection officer, and an external independent advisor with lived experience of algorithmic bias. The panel meets monthly and has a formal power to delay any AI deployment pending review.

A tech startup replaced their informal 'engineering decides' culture with a mandatory 48-hour cross-functional review for any AI feature affecting users directly. The process added two days to their sprint cycle and prevented three features they later agreed would have caused real harm.

A university included student representatives in its AI governance committee, a decision that surfaced concerns about AI-assisted plagiarism detection creating false positives for non-native English speakers, something staff reviewers had not flagged.

"Diversity in this group is not optional, it is how you catch the blind spots."

ACTIONS TO TAKE THIS MONTH

1. Identify 4–6 people across different functions and seniority levels who will form your working group.
2. Include at least one person whose role involves direct community or customer interaction.
3. Define the group's mandate clearly: what decisions require their review, what is their authority?
4. Schedule a standing monthly meeting and document decisions with reasoning.
5. Consider an external advisor (an academic, ethicist, or community advocate) for independent perspective.

5

Require Transparency in Vendor Contracts

YOUR LEVERAGE IS GREATER THAN YOU THINK

Most businesses accept the default terms of AI platforms without negotiation because they assume they lack leverage. But as AI literacy grows and regulatory pressure increases, vendors are increasingly willing to negotiate transparency provisions. The act of asking signals market expectations and, at scale, shifts industry standards.

Negotiate for data transparency clauses in your AI vendor agreements.

Ask vendors to specify what data their models were trained on, how bias is identified and mitigated, and what recourse you have when the tool produces harmful outputs.

WHAT THIS LOOKS LIKE IN PRACTICE

A media company added a 'Model Transparency Rider' to their standard vendor contract template, a one-page addendum requiring vendors to disclose training data sources, bias testing methodology, and incident response obligations. Three of four vendors signed it without resistance.

A healthcare organisation required any AI diagnostic tool vendor to provide documented evidence of performance testing across demographic subgroups (age, ethnicity, gender) before procurement. This is now increasingly standard in regulated healthcare procurement, but it started as a voluntary ask.

A pension fund made their vendor contracts public on request as part of their responsible investment commitments. The transparency requirement prompted their AI payroll vendor to improve their own documentation because the fund's reputation depended on it.

"The act of asking positions you as a sophisticated, accountable buyer and the market will eventually respond to the customer pressure."

ACTIONS TO TAKE THIS MONTH

1. Review your current AI vendor contracts: do any include transparency, bias testing, or incident reporting provisions?
2. Draft a short Model Transparency Rider to include in new contracts and renewals.
3. Ask vendors: What was this model trained on? Has it been tested for demographic bias? What happens when it makes a harmful mistake?
4. Create an internal escalation process for when a vendor cannot or will not answer these questions.

Adopt Wide Metrics for AI Success

WHAT YOU MEASURE IS WHAT YOU MANAGE

Most AI deployments are evaluated on efficiency and cost savings alone. This is like measuring the health of a forest by the volume of timber harvested. Healthy living systems, and healthy organisations, require wide metrics: measures that capture the vitality of the whole, not just the extractable surplus. Regulators, investors, and customers are all moving in this direction.

Stop measuring AI success only through efficiency and cost savings. Add metrics that capture: **employee wellbeing** (does this AI reduce cognitive burden or increase anxiety?), **equity** (does this AI produce different outcomes for different groups?), **customer trust** (are people comfortable with how we use AI?), and **environmental impact**.

WHAT THIS LOOKS LIKE IN PRACTICE

A logistics company added 'driver wellbeing score' and 'route equity index' alongside delivery-time efficiency when evaluating their AI routing platform. They discovered the most efficient routes systematically avoided lower-income postcodes reducing delivery access inequity that the efficiency metric alone would never have revealed.

A software company began measuring 'AI-assisted vs. AI-replaced' task ratios after employees reported feeling deskilled by copilot tools. The metric surfaced a pattern where junior developers were learning less, a long-term risk to team capability that cost savings hadn't captured.

A retail bank added 'customer trust sentiment' (from NPS surveys specifically about AI-driven interactions) and 'approval rate equity by demographic' to their AI KPI dashboard. It was the first time their model risk team and their diversity team had shared a reporting line.

"Measuring success through the vitality of the whole, not just the extractable surplus, is not a luxury, it's where the market is heading."

ACTIONS TO TAKE THIS MONTH

1. List your current AI success metrics. Are they purely operational? Add at least two wellbeing or equity measures.
2. Survey employees specifically about AI tools: do they feel supported or surveilled? Capable or redundant?
3. Track whether AI outputs differ by customer demographic such as age, gender, geography, ethnicity.
4. Include AI's environmental footprint in any sustainability or ESG reporting.
5. Share AI performance metrics with employees who are affected by them, this transparency builds trust.

Create a Detailed Impact Assessment Process

MOST AI HARM IS DETECTABLE BEFORE DEPLOYMENT

The majority of documented AI harms (biased hiring decisions, discriminatory credit scoring, inequitable content moderation. etc.) were theoretically detectable before deployment.

A structured, facilitated impact assessment is not a bureaucratic hurdle, but rather it is a two-hour investment that can prevent years of reputational, legal, and human damage.

Before deploying AI in contexts that affect for example hiring, promotion, customer service, credit decisions, or content generation, run a structured assessment: Who could be harmed by this system? What biases might be embedded in the training data? How will we detect harm if it occurs? What is the process for correction?

WHAT THIS LOOKS LIKE IN PRACTICE

A financial services firm adopted a four-question pre-deployment checklist for any AI system affecting customers: (1) Who is most vulnerable to errors? (2) What historic inequalities might this data encode? (3) How will we know if harm occurs? (4) What is our remediation process? The checklist takes 90 minutes to complete and has become standard practice.

A council ran a community consultation as part of their AI impact assessment before deploying predictive tools in housing allocation. Residents from affected communities identified assumptions in the model's training data that internal reviewers had missed and the tool was modified before go-live.

A global retailer discovered through impact assessment that their AI-generated product descriptions used systematically different language for the same products when targeting different demographic groups subtly reinforcing stereotypes. The assessment prevented launch and prompted a model retraining.

"A two-hour assessment can prevent years of reputational, legal, and human damage. Most AI harm is theoretically detectable before launch."

ACTIONS TO TAKE THIS MONTH

1. Identify which AI use cases in your organisation are 'high stakes' involving for example hiring, credit, health, or community impact.
2. For each high-stakes use case, schedule a 90-minute impact assessment session with your ethics working group.
3. Use a structured template: potential harms, affected groups, bias risks, detection mechanism, remediation path.
4. Document the assessment and make it part of the formal deployment approval process.
5. Review completed assessments 6–12 months post-deployment to see if predicted risks materialised.

8

Invest in AI Literacy Across Your Organisation

ETHICS REQUIRES INFORMED PARTICIPANTS

Ethical AI use is impossible if employees do not understand how AI works, where its limits lie, or how to question its outputs. Widespread AI literacy is not a nice-to-have, but rather it is the precondition for genuine accountability.

Invest in education that goes beyond 'how to use ChatGPT' to include: how AI models work, what their limitations are, where bias comes from, and how to critically evaluate AI outputs.

Employees who understand AI are less likely to over-rely on it, more likely to catch errors, and better equipped to raise concerns when something feels wrong.

WHAT THIS LOOKS LIKE IN PRACTICE

A professional services firm ran a series of 'AI literacy lunches', informal 45-minute sessions exploring topics like 'where does AI training data come from?', 'what is hallucination and why does it happen?', and 'how do I know when to trust an AI output?'. Attendance was voluntary; within three months, 80% of staff had attended at least one.

A hospital trust trained clinical staff not just in how to use AI diagnostic tools but in how to constructively challenge them, including a protocol for escalating when a clinician's professional judgment conflicted with an AI recommendation.

A school integrated 'AI critical thinking' into existing digital literacy lessons teaching students to identify AI-generated content, understand data provenance, and question algorithmic recommendations. They found that staff literacy improved as teachers prepared the lessons.

"An organisation where only IT staff understand AI is one where most of the people affected by it cannot meaningfully participate in its oversight."

ACTIONS TO TAKE THIS MONTH

F Survey employees on their AI literacy. Not to assess them, but to understand where education is most needed.

G Design tiered training: foundational literacy for all staff, deeper technical understanding for those working with AI daily.

H Include 'how to raise AI concerns' in your training, because you want to make escalation pathways explicit and safe.

I Train managers to receive and act on AI concerns from their teams without dismissing them as technical matters.

Í Revisit and update training annually as AI capabilities and risks evolve quickly.

Use AI to Regenerate, Not Just Extract

AI AS A FORCE FOR FLOURISHING

The dominant narrative positions AI as a tool for efficiency, doing more with less, faster, cheaper. But AI can also be a tool for healing, connection, and regeneration. Businesses that explore regenerative AI use cases are building towards a competitive position in an economy that is rapidly placing value on social and ecological contribution, not just throughput.

Look for use cases where AI can actively contribute to wellbeing, whether yours, your community's, and the planet's.

Can AI help you reduce waste in your operations? Improve access to your services for underserved populations? Monitor your environmental impact in real time? Support employee mental health and flexibility?

These applications build genuine value and align your AI use with the direction markets and regulation are heading.

WHAT THIS LOOKS LIKE IN PRACTICE

A food supply chain business used AI to redesign delivery routes to minimise food waste at collection points by combining logistics optimisation with real-time temperature monitoring. The result reduced waste and created a new service offering for smaller food producers who couldn't previously access cold-chain logistics.

A housing association used an AI communication tool to identify tenants showing early signs of financial distress, based on changes in payment patterns, and proactively connected them with support services before arrears escalated. Eviction rates in the pilot area dropped.

A consumer goods company partnered with a university to use AI for real-time monitoring of biodiversity indicators in their agricultural supply chain, going beyond carbon metrics to track soil health, pollinator activity, and water quality. The data informed supplier support programmes and became a differentiating marketing asset.

"AI trained to regenerate rather than extract creates lasting value for your business, your community, and the planet."

ACTIONS TO TAKE THIS MONTH

1. Host a 'regenerative AI ideation' session: what problems in your community, supply chain, or ecosystem could AI help solve?
2. Identify one operational process where AI could reduce waste or environmental harm, not just cost.
3. Explore whether your AI use could be extended to serve customers or communities who currently lack access.
4. Look for ways to share AI-generated insights with suppliers, partners, or communities, not just keep them internal.
5. Measure the social and ecological outcomes of AI use, not just the financial ones.

Join the Movement for Responsible AI Standards

INDIVIDUAL ACTIONS ALONE ARE NOT ENOUGH

Individual business decisions matter, but industry norms, regulatory frameworks, and platform architectures are shaped by collective action. Research on systems change shows that tipping points occur when just 15–25% of participants shift behaviour.

Business leaders who engage in governance conversations now are not late, they are early enough to shape the outcome.

No single company can set industry standards alone, but collectively, businesses have enormous influence over AI platform behavior and regulatory frameworks. Engage with industry coalitions, standards bodies, and policy conversations. Support the development of AI governance frameworks that prioritise Life-alignment, not just risk mitigation. Your voice, as a business leader, as a customer of AI platforms, as a member of civil society, matters more than you may realise.

WHAT THIS LOOKS LIKE IN PRACTICE

A group of mid-market businesses formed a purchasing coalition that collectively wrote to three AI platform providers requesting transparent energy disclosures and bias audit reports as conditions of contract renewal. Two of three providers responded with improved documentation within six months.

An industry association incorporated AI ethics into their annual conference programme and commissioned a sector-specific guidance note on responsible AI deployment, creating shared standards that individual members could reference in vendor negotiations.

A founder joined her country's national AI advisory panel as the sole SME representative, bringing perspectives on practical constraints that large-enterprise voices routinely missed. She described it as the highest-leverage 10 hours she spent all year.

"Tipping points occur when just 15–25% of participants shift behaviour. You do not need to move the whole system, just enough of it."

ACTIONS TO TAKE THIS MONTH

1. Identify two industry coalitions or standards bodies working on AI governance and sign up to their communications.
2. Engage with your industry association: is AI ethics on their agenda? If not, propose it.
3. Respond to public consultations on AI regulation. Regulators genuinely rely on business input.
4. Share what you are learning with your network: peer accountability accelerates adoption.
5. Support organisations developing Life-aligned AI frameworks, financially, through partnership, or through advocacy.

The Train May Have Left the Station But We Are the Engineers

The decisions being made right now about how AI is trained, governed, and deployed will define its trajectory for decades. The window for shaping those decisions is not infinite.

But every purchasing decision you make is a vote for what kind of AI gets built.

Every standard you hold your vendors to is a signal to the market.

Every conversation you have internally about what your AI is actually optimising for is an act of leadership that ripples outward.

The question is not whether AI will transform your industry.

It will.

The question is: **Will you be someone who shaped that transformation toward Life or someone who accepted whatever arrived?**

This guide is drawn from the *Regenerative AI Ethics* framework, co-authored by Martina Doleshal. The framework is grounded in a year-long transdisciplinary study of Living Systems Science, Biomimicry, Indigenous Wisdom Traditions, Regenerative Principles, and Complexity Science.

Key data sourced from:

Capgemini: Developing Sustainable Gen AI Report (2025)

OECD AI Water Consumption Study (2023)

Goldman Sachs Power Demand Analysis (2024)

NPR / Google & Microsoft Emissions Reporting (2024)

MIT Technology Review (2019)

Strubell et al. (2019)

IEEE SA (2023)

For help implementing any of these tips of for development of your AI strategy,

reach out to **Martina Doleshal**,

by emailing: martina@doleshalconsulting.com

and include **[Ethical AI Use]** in the subject of your email.