People First, Always: AI Strategy & Decision Toolkit

Designing AI investment, governance and value creation through people, not just technology

This resource is designed to help leaders place people at the centre of AI investment, governance and value creation.

CONTENTS

This resource pack accompanies the People First, Always webinar. It is designed to surface the right questions for your organisation , not to provide a template answer.

01 Questions To Ask Before You Scale

Take these into your next leadership meeting. Honest answers will tell you where you are.

1	**Can you name the person accountable for each AI decision — not the tool owner, the decision owner?** Not a committee. A named individual who will be held responsible if the system produces a harmful outcome.
2	**Is your re-skilling budget proportionate to your automation ambition?** If your automation investment is growing and your capability development budget is flat, the gap between them is your risk.
3	**What would your workforce say your AI strategy is actually for?** Not what leadership says it is for — what the people delivering the work day-to-day believe it is for.
4	**Have role impacts been formally mapped before any AI goes live?** Before deployment, not discovered after adoption fails. Which roles change, which disappear, which require new capability.
5	**Is your AI investment case built on people assumptions that have been tested — or assumed?** Adoption timelines, capability requirements, management readiness, workforce trust. Were they evidenced before approval, or written in to make the numbers work?

Most AI investment cases are built with precise technology costs and approximate people assumptions. These six questions surface the human variables that determine whether the financial model holds, before the money is committed.

1

What level of adoption is required for this to deliver ROI?

Business outcome: Investment case integrity; the return projection reflects a named, tested adoption assumption rather than optimism.

People dimension: Adoption rate is the variable underneath every financial model. If the target is never named, nobody is accountable for closing the gap between the projection and reality.

2

What behaviours must change for this to succeed?

Business outcome: Sustainable value capture, not just deployment.

People dimension: The specific habits, tasks and processes that must be redesigned around the tool rather than bolted onto existing ones. If these are not named before go-live, they are discovered during the adoption failure.

3

Where could lack of trust or understanding reduce impact?

Business outcome: Adoption depth, not just adoption breadth. A tool can show 80% activation while 60% of users are running shadow processes in parallel.

People dimension: The teams, roles and decision types most likely to generate resistance or workaround behaviour. Trust failures in AI adoption follow a predictable pattern — this question identifies where they will appear before they do.

4

Who is accountable for ensuring people adopt and use this effectively?

Business outcome: Accountability clarity; somebody's performance is assessed on whether the workforce actually uses the tool, not just whether it is available.

People dimension: Adoption is a people-management function, not a technology function. Without a named individual who owns this, adoption becomes everyone's responsibility and therefore nobody's. ***Most AI deployments name a tool owner. Almost none name an adoption owner.***

5

What would failure look like and what would it cost?

Business outcome: Risk-adjusted decision-making. A board that approves an investment without a named failure scenario has not governed the investment.

People dimension: The workforce is both the most common cause of failure and the most effective early warning system. Adoption rates, sentiment and escalation behaviour are the leading indicators of programme failure — visible months before it shows up in financial results.

6

What would we stop doing if this initiative didn't exist?

Business outcome: Strategic alignment; the initiative is funded because it solves a specific problem, not because AI is a board priority.

People dimension: the answer is almost always about human time and attention. What people are being asked to stop doing, deprioritise, or do differently to make capacity for the new programme. If the answer is nothing, the investment case is weaker than it appears.

IF YOU CANNOT ANSWER THESE CONFIDENTLY

> These questions separate organisations that get value from AI from those that generate activity. Seven Palms Consultancy's Executive AI Sprint is designed specifically to work through these with your leadership team before you scale.

02 Governance Health Check: Ten Indicators

Rate each Red, Amber, or Green. **More than three Reds means your governance is performative, not functional.**

	Governance Indicator	
1	**Named individual (not a committee) accountable for each AI deployment.** Not who chairs the AI governance forum. Who is personally responsible if this system produces a harmful outcome.	🔴 Red 🟡 Amber 🟢 Green
2	**Board receives AI performance reporting in non-technical language at least quarterly.** Reporting the board can interrogate — not technical dashboards they nod at.	🔴 Red 🟡 Amber 🟢 Green
3	**Human oversight roles have defined competence, time and authority requirements.** EU AI Act Article 14 requires meaningful oversight not a rubber-stamp process at volume.	🔴 Red 🟡 Amber 🟢 Green
4	**Escalation path exists and has been used in the last 12 months.** If no escalations have ever been raised, the path is not working, not because everything is fine.	🔴 Red 🟡 Amber 🟢 Green
5	**AI use cases are reviewed against original approved purpose at least annually.** AI systems drift. A tool approved for one use gets extended. Someone needs to be checking.	🔴 Red 🟡 Amber 🟢 Green
6	**Workforce has been formally informed of AI systems affecting their work, before deployment.** GDPR Article 22 — Transparency requirement for automated decisions with significant effect.	🔴 Red 🟡 Amber 🟢 Green
7	**Reskilling is budgeted as a line item in the AI programme, not retrospective.** Capability gaps discovered after go-live cost significantly more to close than gaps mapped before.	🔴 Red 🟡 Amber 🟢 Green
8	**A formal AI strategy exists and success metrics were defined before approval.** Enterprises with a formal AI strategy: 80% success rate. Without: 37%. — Writer/Wharton 2025	🔴 Red 🟡 Amber 🟢 Green
9	**Role impacts have been mapped and documented for all material AI deployments.** Before deployment — not discovered during the adoption failure post-mortem.	🔴 Red 🟡 Amber 🟢 Green
10	**Regulatory compliance is treated as a governance framework, not a minimum floor.** EU AI Act · GDPR Article 22 · CSRD. A floor protects you today. A framework protects you tomorrow.	🔴 Red 🟡 Amber 🟢 Green

If you answered Red or Amber on more than three indicators, your governance exists on paper but not in practice. That is the most common and most costly governance failure mode in enterprise AI.

03 The People Metrics Gap

What should be in every AI investment case. Most AI investment cases are built with precise technology costs and approximate people assumptions. The metrics below most reliably predict whether AI investment will deliver value and are most consistently absent from business cases before approval.

Focus on:

Adoption rate
Are people actually using it?

Time-to-competence
How quickly can people use it effectively?

Decision quality
Are decisions improving?

Rework reduction
Are errors and duplication decreasing?

Confidence and trust
Do people trust the outputs?

Process efficiency
Is time meaningfully reduced?

Adoption Readiness

Metric	Why it predicts success	Source
Psychological safety index by team	BCG 2025: When leaders demonstrate strong AI support, positive employee sentiment rises from 15% to 55%	BCG 2025
Shadow AI usage rate	MIT NANDA 2025: Widespread unsanctioned tool use signals trust failure and adoption risk before it becomes visible	MIT NANDA 2025
Manager confidence in AI oversight	Regular AI use: leaders 75%+, frontline 51% and stalling. Gap indicates top-heavy adoption that will not scale	BCG 2025

Capability Gap

Metric	Why it predicts success	Source
Skills gap size vs deployment requirement	Measured before go-live. 70.89% of EU enterprises cite lack of expertise as primary AI barrier	Eurostat 2025
Time available for human oversight per decision	EU AI Act Article 14 — meaningful oversight requires documented time and competence, not just a process step	EU AI Act 2024
Proportion of managers with AI governance training	Governance delegated to technical teams underperforms governance shaped by senior leadership	Deloitte 2026

Leadership Alignment

Metric	Why it predicts success	Source
C-suite sponsorship continuity beyond 6 months	68% success rate with sustained CEO involvement vs 11% when sponsorship lapses — the single biggest predictor of success	RAND / MIT / McKinsey
Executive agreement on success metrics pre-approval	73% of failed AI projects lack clear executive alignment on what success means	Pertama Partners 2025–26
Board AI literacy assessment	Only 34% of organisations truly reimagining business with AI. Most board AI discussions remain ceremonial	Deloitte 2026

Workforce Impact

Metric	Why it predicts success	Source
Role impact mapping completion rate	EU AI Act high-risk classification. Most HR/workforce AI requires conformity assessment including role impact before deployment	EU AI Act 2024
Internal mobility pathway availability	WEF Jan 2026 — 92M jobs displaced by 2030. Internal mobility structures reduce redundancy cost and workforce anxiety	WEF 2026
AI skills wage premium gap within the organisation	Workers with AI skills earn 56% more. Internal gap predicts both retention risk and capability shortfall	PwC 2025

Assess Before You Scale: Technology can be deployed quickly, but readiness cannot.

Area	Score (1–5)	Comments
Leadership alignment		
Workforce readiness		
Adoption capability		
Governance maturity		
Change capacity		

Low scores indicate risk to ROI - Regardless of technical capability.

04 Regulatory Snapshot

What leaders need to know right now.

Three regulatory frameworks are directly shaping the financial and governance risk of AI investment decisions today. The question for each is not whether you are compliant, it is whether your governance is built to remain compliant as these frameworks mature.

THIS IS CONTEXTUAL AWARENESS, NOT LEGAL ADVICE.

EU AI Act *(Phased implementation 2024–26)*

Risk-tiered obligations. Most HR and workforce AI [Hiring screening, performance management, workforce planning] will be classified high-risk. High-risk AI requires conformity assessment, human oversight documentation and transparency before deployment. **Fines up to €35M or 7% of global turnover.**

The board question: "Which of our AI deployments would be classified high-risk and have we completed the required conformity assessments?"

EU AI Act 2024. Gartner: 40%+ of agentic AI projects predicted cancelled by 2027.

GDPR / UK GDPR — Article 22

Individuals retain the legal right to human review of any decision made solely by automated processing that has legal or significant effect. Applies to AI-assisted hiring, performance management, and credit decisions. 'Human in the loop' in a process document does not satisfy Article 22 "The human must have genuine competence, time and authority to intervene".

The board question: "For each AI system influencing decisions about people can we demonstrate that human oversight is meaningful, not perfunctory?"

GDPR Article 22. ICO guidance on automated decision-making, UK.

CSRD *(Corporate Sustainability Reporting Directive)*

Workforce impact of AI is increasingly a reportable ESG matter for large organisations. CSRD requires disclosure of how technology adoption affects the workforce; including displacement, reskilling investment and inclusion.

The board question: "Is our AI programme's workforce impact [role changes, displacement risk, re-skilling investment] captured in our sustainability reporting?"

CSRD Directive 2022/2464/EU. WEF Future of Jobs Jan 2026.

THE GOVERNANCE QUESTION THAT SITS ABOVE ALL THREE

Are you treating regulatory compliance as a floor (the minimum to avoid a fine) or as a governance framework — the structure that keeps your AI investment aligned to your values as the landscape shifts?

Organisations that treat it as a floor face continuous retrofit cost. Those that treat it as a framework build once and adapt incrementally.

Note: This is a contextual overview for leadership awareness, not legal advice. Seek qualified legal counsel for compliance decisions specific to your organisation.

05 First 90 Days of Responsible AI Adoption

A framework for leaders starting, or restarting, an AI programme with people at the centre.

The organisations that get value from AI do not move faster than others. They move more deliberately. They diagnose before they design, align before they build and pilot with purpose. These ninety days are about building the foundations that make scale possible and sustainable.

1 2 3

Days 1–30

Diagnose before you design

- Map current AI activity [sanctioned and unsanctioned]
- Confirm executive intent is shared, not assumed
- Assess workforce sentiment and anxiety levels
- Identify who makes decisions AI might affect
- Name the accountability owner — not the tool owner

Output: Decision accountability map · Shadow AI audit · Workforce sentiment baseline

Days 31–60

Align before you build

- Establish executive intent
- Assess reskilling needs and gaps
- Validate governance against live use cases
- Document decision frameworks
- Define which decisions remain human-led
- Formalise governance structures
- Begin capability mapping: gaps and who closes them

Output: Agreed AI strategy statement · Human oversight standards · Capability gap assessment

Days 61–90

Pilot with purpose

- Communicate progress to the organisation
- Launch reskilling programmes
- Run one bounded use case with measurable people and performance metrics
- Build in review gates, not just launch dates
- Document what you learn, not just what worked
- Report to the board on human impact, not just efficiency gains

Output: Pilot review report with people metrics · Board-level impact statement · Go/no-go for scale

2x

Financial returns

More likely to report significant financial returns when workflows are redesigned before tools are selected - McKinsey 2025

68%

vs 11%

AI success rate: sustained CEO involvement vs sponsorship lost within 6 months — RAND / MIT / McKinsey - 2,400+ initiatives

88%

PoCs fail to scale

of AI proof-of-concepts never reach production; capability and governance gaps are primary causes - IDC

Common People-Related Failure Patterns

AI initiatives often lose value due to:

- Technology moving ahead of workforce readiness
- No ownership of behavioural adoption
- Governance focused on systems, not decisions and accountability
- Training delivered, but capability not embedded
- Activity measured, but human impact ignored
- Leadership intent not reflected in day-to-day behaviour

> When people are not designed into the solution, value is delayed, diluted or lost.

Sources: McKinsey 2025 · RAND Corporation 2024 · MIT NANDA GenAI Divide 2025 · IDC · Pertama Partners 2025–26

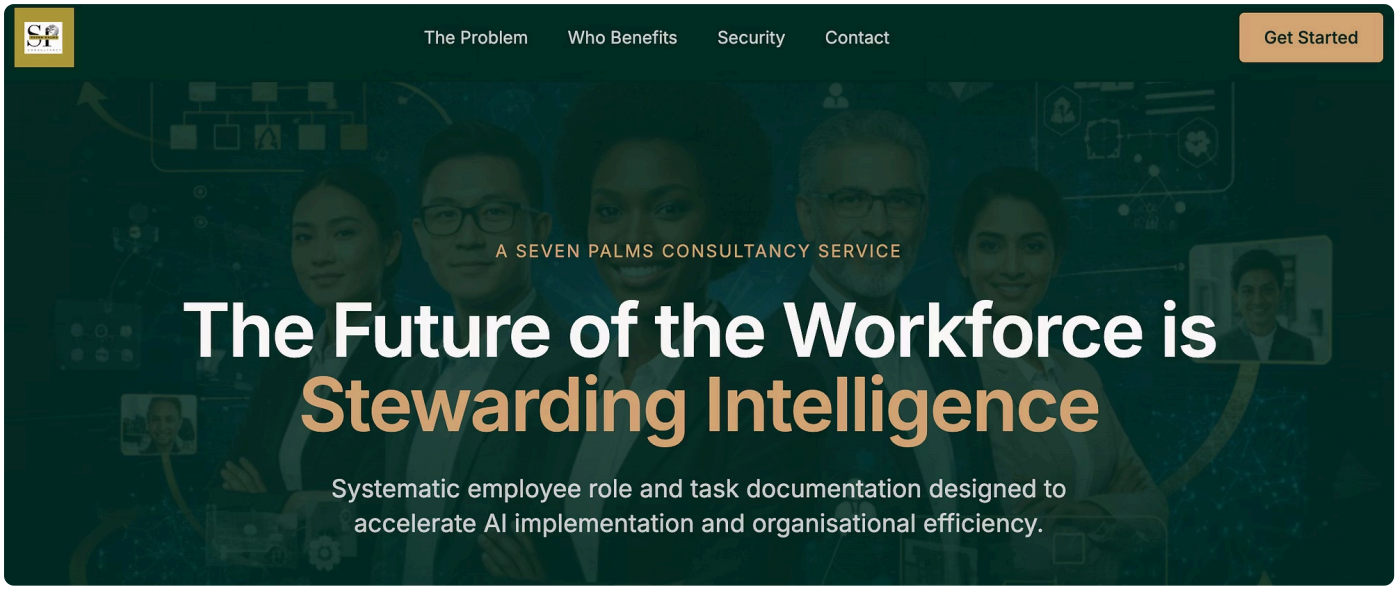

THE CONNECTION TO STEWARDING INTELLIGENCE™ : The 90-day framework creates the conditions for AI adoption. Stewarding Intelligence™creates the organisational knowledge that makes those conditions work in practice; documenting the roles, tasks, decision-making authority, and workflows that AI systems need to be effective from day one. The two work together: the framework tells you how to approach adoption responsibly. Stewarding Intelligence™ gives your AI systems and your people the clarity they need to succeed.

06 Stewarding Intelligence

A Seven Palms Consultancy service and the practical mechanism that makes people-first AI adoption work.

One of the most consistent findings across AI adoption research is that organisations spend **40–60 hours per role** re-teaching AI systems the organisational context they need to be effective. Every new tool. Every new deployment. The same institutional knowledge re-entered from scratch.

Stewarding Intelligence™ solves this. It is a systematic role and task documentation service that creates comprehensive, privacy-protected maps of how your organisation actually works; the responsibilities, workflows, decision-making authority, and institutional knowledge that currently exists only in people's heads.

The Problem It Solves

- **The AI context gap**
  LLMs require detailed organisational context to be effective. Without it, AI outputs miss the mark, require constant correction, and never deliver the productivity gains the investment case promised.
- **The organisational knowledge gap**
  Role clarity is often informal and undocumented. Institutional knowledge leaves when employees depart. Onboarding takes 3–6 months because nobody has written down how things actually work.

What It Delivers

- 40–60% faster AI tool implementation - context built in, not retrofitted
- Reduced onboarding from months to weeks - documented workflows from day one
- Privacy-protected by design - role documentation, not individual surveillance
- Living documentation - updated as roles evolve, not a one-time project
- Employee-centric - every participant reviews and approves their documentation before finalisation

BUILT ON THE PRINCIPLES THAT UNDERPIN THIS SESSION;
People first means giving your AI systems the human context they need to work well — and giving your people the clarity they need to work alongside AI without anxiety or ambiguity.

Stewarding Intelligence™ is the operational expression of everything discussed in this session: documented accountability, mapped roles, tested assumptions, and institutional knowledge that stays when people leave.

Executive AI Sprint

A structured leadership session using the Stewarding Intelligence™ framework, aligning your executive team on AI strategy, governance, and people readiness before you scale.

Stewarding Intelligence™

The documentation service that gives your AI systems and your people the organisational context they need from day one. Privacy-protected. Employee-approved. Built to scale.

Start a Conversation
info@sevenpalmsgroup.com
+44 (0)735 621 9967
sevenpalms.group

Before you invest further...Ask these critical questions...

- Are our people ready - not just our systems?
- Do we understand the behaviours required for success?
- Are we measuring adoption - not just output?
- Are decision ownership and accountability clear?
- Are we prepared to stop if value is not realised?

If these questions resonate or you would like to explore how this applies within your organisation - Seven Palms Consultancy supports organisations in designing people-first AI and digital strategies that deliver measurable value. Let's continue the conversation.

> *PEOPLE FIRST, ALWAYS · Seven Palms Consultancy · March 2026*

07 Governance Reflection — Session Reference

Key questions from the session to carry back to your organisation. These questions were presented during the session as reflection prompts. They are reproduced here as a reference to use with your own leadership team. Each has a rationale, the reason it matters and what an honest answer tells you.

1 What governance questions ensure AI remains aligned to organisational priorities?

Governance that governs technical performance but not strategic alignment is governance theatre. The question is not whether your AI works — it is whether it is still doing what you said it would do, for the reasons you said it would do it. *Gartner 2026 · Deloitte State of AI 2026*

2 Is executive intent aligned?

Can every member of your leadership team state your AI strategy in one sentence — and say the same thing? Misalignment at the top is the most common and most expensive source of failed AI investment. *73% of failed projects lack clear executive alignment — RAND/MIT/McKinsey synthesis*

3 Is decision accountability named?

Not which committee owns AI governance. Which individual is personally accountable for each decision the system influences and will be held responsible if it produces a harmful outcome. *RAND Corporation 2024*

4 Have role impacts been mapped?

Before deployment, not after adoption fails. Which roles change, which disappear, which require new capability and who is accountable for that transition. *EU AI Act 2024*

5 Is re-skilling funded and prioritised?

As a capital investment proportionate to your automation ambition — not a training budget line added after the technology spend is approved. *Deloitte State of AI 2026 · Eurostat 2025*

6 What is the cost assumption in your AI investment case if this programme fails - and is that figure in the business case?

7 Is your AI investment case built on people assumptions that have been tested - or assumed?

Each reflection question maps to a primary cost pillar and often touches a second. The mapping shows which failure mode each question is designed to prevent.

#	Question · Rationale	Primary Pillar	Also Touches
1	**What should leaders formally document before scaling AI?** Documentation creates the audit trail and accountability structure. Without it, governance is unverifiable. The absence of documentation is also a direct cause of rework - decisions made without a record get made again, differently, later.	**Governance**	Rework; Undocumented decisions are re-made, at cost
2	**What governance questions ensure AI remains aligned to organisational priorities?** Alignment to priorities is a governance function, it requires a mechanism to check drift. But it only works if leadership has named the priorities clearly and owns the alignment question.	**Governance**	Leadership; Leadership must name the priorities governance checks against
3	**Is executive intent aligned?** Misalignment at the top is the most common and most expensive source of failed AI investment. It cannot be assumed it must be tested. When leaders say different things, the workforce receives contradictory signals and adoption stalls.	**Leadership**	Resistance; Contradictory signals from leaders are a direct adoption barrier
4	**Is decision accountability named?** Accountability is personal before it is structural. This question points at a leadership failure, nobody has named who owns the consequence. Governance provides the framework, but leadership has to fill it with actual names.	**Leadership**	Governance; Governance without named individuals is governance theatre
5	**Have role impacts been mapped?** Unmapped role impacts are the single biggest driver of workforce anxiety and adoption resistance. People resist what they cannot see coming. The EU AI Act also requires documented impact assessment for high-risk AI before deployment.	**Resistance**	Governance; EU AI Act requires documented impact assessment pre-deployment
6	**Is re-skilling funded and prioritised?** Unfunded re-skilling is the structural cause of adoption failure. Capability gaps discovered after go-live cost significantly more to close than gaps mapped before. If reskilling is not budgeted, the organisation will pay for it anyway in rework, resistance, and failed deployment.	**Resistance**	Rework; Underfunded capability = rework cost after go-live
7	**What is the cost assumption in your AI investment case if this programme fails and is that figure in the business case?** Most AI investment cases contain an ROI projection but no failure cost assumption. If the business case does not model what happens when it goes wrong, the board has approved a one-sided financial argument. $4.2M is the average sunk cost per abandoned AI initiative. That figure belongs in every business case before approval, not in the post-mortem.	**Rework**	Governance; Absent failure cost = unvetted investment approval

Each reflection question maps to a primary cost pillar and often touches a second. The mapping shows which failure mode each question is designed to prevent.

#	Question · Rationale	Primary Pillar	Also Touches
8	**Who reviewed the last AI audit report — and what did they do with it?** The absence of a documented response is evidence that oversight is performative. If nobody acted on the audit, the governance structure produced a report but not governance. A governance process that generates findings without generating action is a liability, not a safeguard.	**Governance**	Leadership; Unacted audit findings = leadership accountability failure
9	**Are we measuring what we said we would measure?** Most AI investment cases contain success metrics agreed at approval. The question - asked at six, twelve, and eighteen months post-deployment - is whether those metrics are still being tracked, whether the methodology has held, and whether results are being reported to the people who approved the investment. In most organisations, investment case metrics are quietly abandoned within three months of go-live and replaced with whatever is easy to report.	**Rework**	Governance; Abandoned metrics hide rework cost until it is too late
10	**Does the board understand what it has actually approved?** Not commercially, technically. Do board members know which decisions the AI system makes or influences? Which populations it affects? What happens when it is wrong? Most boards do not, because nobody translated the technical reality into board-level language. A board that cannot interrogate these questions has not governed the investment, it has endorsed it.	**Leadership**	Governance; Boards cannot govern what they do not understand
11	**Who in your organisation owns the accountability for AI decisions - not the tool, the decision?** Accountability is personal before it is structural. Most organisations name the tool owner, the vendor, or the team, not the individual accountable for the decision the system influences. A named tool owner and a named decision owner are different people with different responsibilities.	**Leadership**	Governance; Named tool ownership is not decision accountability
12	**Is your AI investment case built on people assumptions that have been tested or assumed?** Most AI investment cases contain untested assumptions: adoption timelines, capability requirements, management readiness, workforce trust, governance capacity. These are the assumptions most likely to be wrong and the last to be challenged before approval. When wrong, they are discovered after go-live when adoption stalls and the rework bill arrives.	**Rework**	Resistance; Untested people assumptions become rework cost
13	**Does your board distinguish between AI as strategic capability versus active experimentation?** Strategic capability and experimentation require different governance, different accountability, and different success metrics. A board that cannot make this distinction will apply the same oversight to a production HR system as to a proof-of-concept and will either over-govern experiments or under-govern live systems.	**Leadership**	Governance; Undifferentiated oversight misapplies governance at both ends

VALUES & SUSTAINABILITY · Questions 14–16 · An extra pillar for consideration, outside the cost framework

These questions do not map to a cost pillar. They sit in a separate frame, not about what AI failure costs, but about whether AI investment is coherent with what the organisation claims to stand for.

14. Do our AI choices reflect the values we claim to hold?

Primary: Values & Sustainability

This is not a compliance question. It is a coherence question. An organisation that publishes a commitment to fairness, transparency, or employee dignity and deploys AI systems that contradict those commitments has a values integrity problem that no governance framework resolves. The question is most powerful when asked by someone who already suspects the answer is no.

15. What is the environmental and infrastructure cost of how we use AI?

Primary: Values & Sustainability

Large-scale AI deployment carries measurable energy and infrastructure cost that most organisations do not account for at investment approval. CSRD reporting obligations are beginning to require disclosure of this. The question is both an ESG accountability question and a cost transparency question, two dimensions that will increasingly overlap as reporting requirements mature.

16. Are we optimising for speed or for long-term responsibility?

Primary: Values & Sustainability

This is the meta-question that sits above all sixteen. Speed and responsibility are not always in conflict, but when they are, the choice an organisation makes reveals its actual priorities, regardless of what its values statement says. It is the question most likely to make a leadership team uncomfortable, and for that reason the most valuable one to ask.

READY TO TAKE THIS FURTHER?

If working through these questions surfaces more gaps than you expected, that is the most common experience of organisations at the start of a serious AI governance conversation. It is not a reason to slow down, it is a reason to build the right foundations before you scale.

Seven Palms Consultancy works with leadership teams to move from honest self-assessment to practical governance that functions in reality, not just on paper. Start with a conversation: info@sevenpalmsgroup.com

08 Recommended Resources - Part One

All data cited in event slides, resource pack and supporting documents.

Primary Annual Reports

Stanford HAI AI Index Report
Annual report published each April. Primary source for global AI investment figures, legislative tracking across 75+ countries, and adoption data. 2025 edition used throughout.
hai.stanford.edu/ai-index · Published annually · April

McKinsey Global Survey on the State of AI
Annual survey published each autumn. Primary source for enterprise adoption rates, leadership factors, and EBIT impact data. Consistent methodology from 2020 enables direct year-on-year comparison.
mckinsey.com · Published annually · Autumn

Investment & Scale

Source	Key Finding	Date
Gartner — Worldwide AI Spending Forecast	$2.52T global AI spend projected for 2026 — a 44% year-on-year increase from $1.5T in 2025. gartner.com	January 2026
Goldman Sachs — AI Infrastructure Report	$527B hyperscaler capex forecast for 2026. goldmansachs.com	December 2025
OECD — AI Investment Data	61% of all global VC went to AI in 2025. oecd.org	February 2026
Pertama Partners — AI Project Failure Statistics 2026	$4.2M average sunk cost per abandoned initiative. $547B of $684B invested in AI in 2025 failed to deliver value. 80.3% overall failure rate (RAND synthesis). 84% of failures are leadership-driven. pertamapartners.com	2025–26 · Point-in-time

Adoption & Workforce

Source	Key Finding	Date
World Economic Forum — Future of Jobs Report 2025	Net 78M new jobs by 2030. 170M created, 92M displaced. 41% of employers plan to reduce headcount within 5 years due to AI. weforum.org	January 2026
BCG — AI at Work 2025	25% of frontline employees receive strong leadership support for AI adoption. Positive sentiment rises from 15% to 55% with strong leadership. bcg.com	July 2025
PwC — AI Jobs Barometer 2025	56% wage premium for workers with AI skills. Skills changing 66% faster in AI-exposed occupations. pwc.com	2025
Eurostat — ICT Usage in Enterprises 2025	20.2% of EU enterprises using AI. 71% cite skills gap as the primary adoption barrier. ec.europa.eu/eurostat	2025
S&P Global Market Intelligence — AI Adoption Survey 2025	42% of companies abandoned most AI initiatives in 2025, up from 17% in 2024. spglobal.com	2025
Layoffs.fyi — Tech Layoff Tracker	55,775 tech jobs lost in 2026 to date as of 14 March. 764 per day. Meta: 16,000 jobs cut alongside $135B AI spend. layoffs.fyi	March 2026 — Continuously updated

Failure, ROI & Readiness

Source	Key Finding	Date
Deloitte — State of AI in the Enterprise 2026	Only 34% truly reimagining business with AI. 42% believe strategy 'highly prepared' but feel less prepared on infrastructure, data, risk, and talent. 66% struggle to establish ROI metrics. deloitte.com	2026
RAND Corporation — AI Project Outcomes Research	68% success rate with sustained CEO involvement vs 11% when sponsorship lapses within 6 months. 73% of failed projects lack clear executive alignment on success metrics. 2,400+ initiatives analysed. rand.org	2024–25
IDC — AI Proof-of-Concept Research	88% of AI proof-of-concepts never reach production. idc.com	Ongoing
MIT NANDA GenAI Divide Report 2025	95% of GenAI pilots fail to scale to measurable P&L impact. Methodology has been contested — used as directional data only, not a primary citation. mitsloan.mit.edu	2025 — Label as directional
Writer / Wharton — State of AI Report 2025	42% of C-suite say AI adoption is tearing their company apart. 80% success with formal AI strategy vs 37% without. writer.com	2025
Gartner — Agentic AI Forecast 2026	40%+ of agentic AI projects predicted to be cancelled by 2027 due to inadequate risk controls. gartner.com	2026
BCG — Digital Transformation Research (850+ companies)	Only 35% of digital transformation initiatives reach stated goals. 70% failure rate consistently identified across editions. bcg.com	2020–21

Governance & Regulation

Source	Key Finding	Date
EU AI Act	Maximum fines €35M or 7% of global annual turnover. Phased enforcement from February 2025. Most HR and workforce AI classified as high-risk. Exceeds GDPR penalties. eur-lex.europa.eu	Official Journal of the EU — 2024
GDPR Article 22	Meaningful human oversight of automated decision-making affecting individuals. Maximum fines €20M or 4% of global annual turnover. gdpr-info.eu	In force since May 2018
CMS Law — GDPR Enforcement Tracker	Cumulative GDPR penalties: €5.88 billion to date. Comprehensive database of all EU enforcement actions. enforcementtracker.com	Continuously updated
CSRD — Corporate Sustainability Reporting Directive	Requires disclosure of AI's workforce impact as part of ESG reporting obligations. eur-lex.europa.eu	In force 2024
Harvard Law School Forum / NACD — Board Oversight of AI	NACD recommends board oversight of AI 'with the same scrutiny as financial risk.' No current express legal obligation to mitigate AI workforce displacement beyond financial disclosure rules. corpgov.law.harvard.edu	2025

Historical Baselines

Source	Key Finding	Date
McKinsey / University of Oxford — Large IT Projects (5,400 projects)	Large IT projects average 45% over budget, delivering 56% less value. $66B total cost overrun across sample. Used as 10-year technology baseline. mckinsey.com	Published 2012 - data through 2015
IBM — Poor Data Quality Cost Estimate	Poor data quality costs US businesses $3.1 trillion annually. Used as 2015–16 governance cost baseline. ibm.com	2016
Harvey Nash / KPMG — CIO Survey	Only 41% of companies had enterprise-wide digital strategy in 2016–17. Only 18% rated digital technology use as effective. kpmg.com	2016–17

Recommended Resources - Part Two

Curated reading, reports and videos.

Reports

The State of Enterprise AI

OpenAI · 2025–26
Real-world data on how organisations are getting value from AI in practice.

The 2026 State of AI Agents

Anthropic · 2026
How enterprises are building and deploying AI agents in 2026. Essential context for the governance of agentic systems, the fastest-growing and least-governed deployment type.

The State of AI Report

McKinsey · 2025
The annual McKinsey State of AI report.

Leading in the Age of AI Agents

BCG · 2025–26
How leaders must navigate a new era defined by AI agents. Covers the leadership and cultural dimensions.

Annual Reports — Follow These Every Year

Stanford HAI AI Index

Stanford University · Published each April · Free
The single most credible annual source on AI investment, governance, adoption, and workforce impact. If you read one report every year on AI, this is it.
hai.stanford.edu/ai-index

WEF Future of Jobs Report

World Economic Forum · Published biennially · Free
The most credible source for workforce displacement and creation projections. Essential context for any board conversation about AI and the workforce.
weforum.org

Deloitte State of AI in the Enterprise

Deloitte · Published annually · Free
Tracks organisational readiness, governance maturity, and the gap between strategic ambition and operational reality. Most useful for the 'how prepared are we really?' conversation.
deloitte.com

Governance & Regulation — Essential Reading

EU AI Act — Plain English Summary

European Commission AI Office · Free
The official summary of the EU AI Act requirements by risk tier. Shorter than most expect. Essential for any organisation operating in or serving the EU.
digital-strategy.ec.europa.eu

NIST AI Risk Management Framework (AI RMF 1.0)

US National Institute of Standards and Technology · Free
The most widely adopted voluntary AI governance framework globally. Practical, non-prescriptive, and free. A useful starting structure — not US-specific despite the source.
nist.gov/artificial-intelligence

IAPP AI Governance Centre — Resource Library

International Association of Privacy Professionals · Free access
The most comprehensive free resource library on AI governance practice, regulation, and accountability. Updated continuously. Includes the AIIGP certification pathway for professionals.
iapp.org/ai-governance

Videos — Curated for time poor executives

Selected for clarity and relevance to people-first strategy, not for technical depth.

Why AI Is Our Ultimate Test and Greatest Invitation

Tristan Harris — TED2025 · ~20 mins
The most important talk for the Values & Sustainability conversation. Harris argues that AI is humanity's ultimate test of wisdom and that we are repeating the same avoidable mistakes made with social media. Urgent, accessible, and not technically dense.
[ted.com — search 'Tristan Harris TED 2025']

Machine Intelligence Makes Human Morals More Important

Zeynep Tufekci — TED · ~18 mins
Harvard professor on why perceived AI objectivity is dangerous and why human oversight is not optional. Directly relevant to the 'does your board understand what it has approved?' governance question.
[ted.com — search 'Tufekci machine intelligence']

3 Principles for Creating Safer AI

Stuart Russell — TED · ~17 mins
UC Berkeley AI scientist on the problem of misaligned objectives and why human oversight needs to be designed in, not bolted on. Useful foundation for governance conversations.
[ted.com — search 'Stuart Russell safer AI']

The Danger of AI is Weirder Than You Think

Janelle Shane — TED · ~12 mins
The most accessible and entertaining talk on AI failure modes and unintended consequences. Recommended for board-level audiences who are new to AI. Explains why human judgment remains essential.
[ted.com — search 'Janelle Shane danger of AI']

What Happens When Our Computers Get Smarter Than We Are?

Nick Bostrom — TED · ~16 mins
Oxford professor on the nature of superintelligence and what 'human-compatible AI' means in practice. The most intellectually rigorous short talk for audiences who want the philosophical foundations of governance.
[ted.com — search 'Nick Bostrom computers smarter']

PEOPLE FIRST, ALWAYS · Seven Palms Consultancy · March 2026
Seven Palms Consultancy supports organisations in designing people-first AI and digital strategies that deliver measurable value. Let's continue the conversation:
info@sevenpalmsgroup.com · sevenpalms.group