

FAB CAMPAIGNS ONE-PAGER · AI GOVERNANCE · VERSION 1.0

10 Questions Every Board Should Ask About AI Governance

Qualifying the AI governance journey of the board you sit on, or advise.

PREPARED BY	Fabrizio De Liberali, Fab Campaigns Ltd	DATE	May 2026
DOCUMENT REF.	FC-2026-AIG-01	VERSION	1.0 — Live
CLASSIFICATION	C1 · Public	PAIRS WITH	FC-2026-CYB-01 · Cyber + Data Risk

Most boards are being asked: “How will AI transform our business?” The more urgent question is the reverse: **what is AI already doing inside our business, and who is accountable for it?** These ten questions are designed to surface the answers, and the silences.

01 • VISIBILITY & INVENTORY

1 What is our complete inventory of AI systems in active use, including AI embedded within the third-party SaaS we already pay for, and who is accountable for keeping it current?

WHY IT MATTERS You cannot govern what you cannot see. Most AI policies sit on top of systems management does not know it owns.

2 Which of those AI systems make or materially influence decisions about customers, employees, or our strategic position, and what risk classification have we assigned to each?

WHY IT MATTERS Not all AI carries equal consequence. A marketing assistant is not the same as an algorithm that ranks CVs, prices risk, or screens transactions.

02 • STRATEGY & RED LINES

3 What is our stated AI risk appetite, and where have we drawn the red lines we will not cross, even if a competitor does?

WHY IT MATTERS Competitive pressure ratchets AI behaviour downward. Without explicit red lines, drift is automatic and invisible.

4 How do we test whether our AI investments are producing the outcomes we claimed, and what are the criteria under which we would stop them?

WHY IT MATTERS AI pilots develop a cultural immunity to being killed. Boards must demand kill criteria before approving spend, not after.

03 • ACCOUNTABILITY & EXPLAINABILITY

5 For every AI system making decisions about people, can we explain how that decision was made, in plain language, at the moment it is challenged?

WHY IT MATTERS *"Because the model said so" fails GDPR Article 22, the EU AI Act, and basic duty of care. Explainability is operational, not documentary.*

6 Can we prove the training data behind our AI, especially third-party models we procure, was lawfully obtained and appropriate for our use case?

WHY IT MATTERS *Data lineage is the new supply-chain risk. Unlawful or biased training data is a latent liability that compounds with every decision made.*

04 • RESILIENCE & DEPENDENCY

7 When an AI system produces a wrong, biased, or harmful output, not if, who detects it, who owns the response, and how quickly can we intervene?

WHY IT MATTERS *AI failures rarely announce themselves. The detection gap is where reputational damage accumulates quietly.*

8 Have we tested what happens if a critical AI vendor withdraws a model, changes its terms, or suffers an outage, and do we have a viable fallback?

WHY IT MATTERS *AI vendor concentration is now a strategic dependency. Without a tested exit, you are subject to their pricing, policy, and geopolitics.*

05 • LITERACY & RIGHTS

9 What AI literacy do our board, executives, and frontline workforce actually have, and how is it measured, not merely asserted?

WHY IT MATTERS *Under the EU AI Act, AI literacy is a legal obligation. Boards cannot oversee what they do not understand.*

10 What rights do our customers and employees have regarding AI decisions about them, and how do we honour those rights in practice, not just in policy?

WHY IT MATTERS *Rights that exist only on paper create regulatory, reputational, and litigation exposure. Real governance is operational.*

THE ETHICAL RATCHET PRINCIPLE

In AI governance, ethics should only tighten, never loosen. Once you cross a red line, you cannot uncross it. The board that draws its red lines before the pressure arrives is the only board that keeps them.

A LIVE DOCUMENT

This is Version 1. I am testing and refining these questions with boards across the UK and Italy through 2026. If you use them, or disagree with any of them, I want to hear from you. **Reputations are built by being cited, not by being followed.**

Fabrizio De Liberali Partnership Architect, Fab Campaigns Ltd Pairs with: 10 Questions on Cyber Governance + Data Risk
