InfoSecAI

A BRIEF FOR BOARDS & AUDIT COMMITTEES

BOARD PACK · 2026

The AI Governance Board Pack

Seven decisions every UK board, audit committee chair and CISO should make before August 2026. A practitioner's brief on the EU AI Act, ISO/IEC 42001, and the new wave of AI scrutiny.

AUTHORED BY

Paul Jolliffe

Founder & Director, InfoSecAI · Senior CISO / vCISO · CISSP · ISO 27001 Lead Auditor · MBA

01 · WHY THIS MATTERS NOW

The board's AI window is closing.

In August 2026, the EU AI Act's general-purpose AI obligations and high-risk system rules begin to bite. ISO/IEC 42001 published in December 2023 is now being asked for in customer assurance questionnaires. The UK ICO has signalled that automated decision-making and AI in HR, credit and education will be priority enforcement areas. **Boards that have not yet made the structural decisions on AI governance are running out of clean air.**

This brief sets out the seven decisions I take every UK board through. Each one has a defined owner, a defined artefact, and a defined date. Together they form a defensible AI governance position: one your auditors, regulators, customers and insurers can see, your CISO and DPO can operate, and your CEO can talk to your shareholders about.

7	90 days	Aug 2026
decisions to make	realistic stand-up time	EU AI Act high-risk gate

> **Reading note.** If you are a board member or audit committee chair, the seven headline decisions and the maturity model are the priority pages. If you are a CISO or DPO operationalising this, the use-case decision tree and common-mistakes pages will save you the most time.

One slide, seven decisions, one accountable line for each.

The board does not need to become an AI governance committee. It needs to make seven decisions and verify that each has an owner, an artefact, and a review cadence.

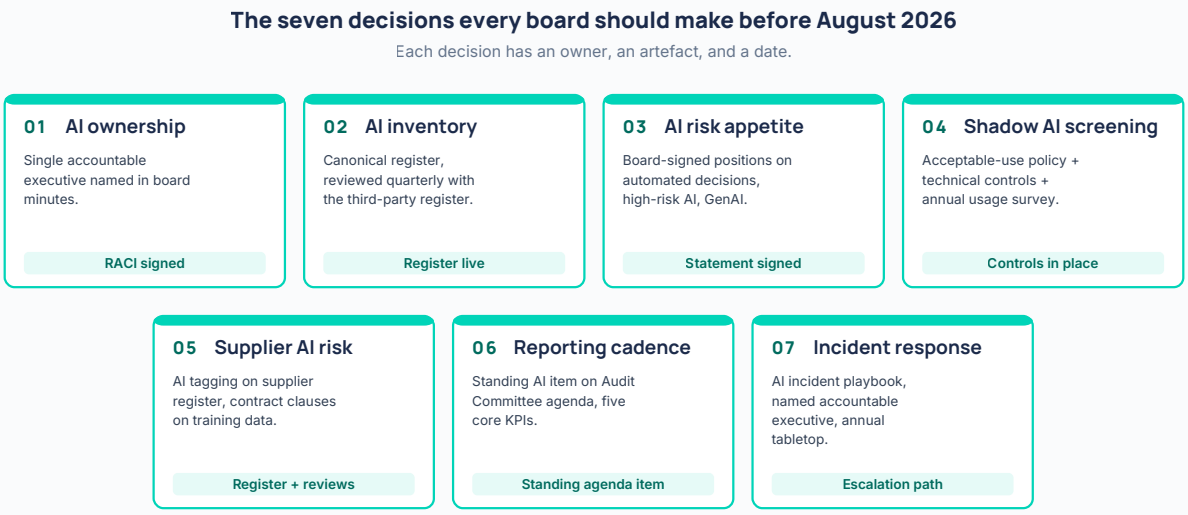

The board's AI governance decision wheel — each decision has an owner and a review cadence.

Decision 01 · Who owns AI governance?

Every framework — ISO 42001, NIST AI RMF, the EU AI Act — points to the same starting move: name an accountable executive. Most UK boards I assess do not have a clean answer. The CIO and the DPO point at each other; the CISO is too often expected to absorb it without explicit delegation.

What good looks like

- Single accountable executive named in the board minutes (typically CIO, CTO, or COO; sometimes CISO).
- RACI signed by Audit Committee covering policy, inventory, risk, supplier review, incident response and reporting.
- Standing item on Audit Committee agenda quarterly.

Decision 02 · What is our AI inventory?

You cannot govern what you cannot see. The AI inventory is the asset register of AI use, and it is the single most under-built artefact in UK organisations right now.

What good looks like

- One register, owned by the AI accountable executive, refreshed at least quarterly.
- Captures: use case, business owner, model / supplier, data classes, decision impact, status, last review.
- Reviewed by Audit Committee at the same cadence as the third-party register.

Decision 03 · What is our risk appetite for AI?

The board owns risk appetite. It cannot delegate this. Most appetite statements I see treat AI as part of "technology risk" — a bucket that is now too coarse. The board needs an explicit position on customer-impacting AI, on automated decision-making, on AI-assisted diligence, and on use of generative AI for regulated communication.

What good looks like

- One-page AI risk appetite statement signed by the board, reviewed annually.
- Explicit positions on: automated decisions affecting customers; high-risk AI use cases under Annex III; cross-border AI deployment; and incidents.
- Translates into clear "yes / no / conditional" for use-case approvals.

Decision 04 · How do we screen for shadow AI?

Shadow AI is the present-tense problem. Employees are using ChatGPT, Claude, Copilot and a hundred other tools on customer data, on confidential pipeline, on internal IP. Most boards I brief have not yet asked the basic question: is this happening, and what controls are in place?

What good looks like

- Acceptable-use policy explicitly covering generative AI, with named approved tools and prohibited actions.
- Network / DLP / browser controls applied to public AI endpoints; logging where applicable.
- Survey of staff usage at least once per year, with named control owner remediating gaps.

Decision 05 · What is our supplier AI risk position?

The largest AI risks most organisations carry are not in their own AI; they are in their suppliers' AI. The cloud productivity stack, the CRM, the recruitment platform, the contract management system — all are quietly adding AI features. Each is a potential data-flow event, a potential automated-decision event, a potential transparency obligation.

What good looks like

- Supplier register tagged for AI use; concentration on key providers (Microsoft, Google, Salesforce) tracked.
- AI-specific clauses in contract renewals: training data, sub-processors, audit rights, model dependency, deletion rights.
- Annual re-attestation, with the supplier register reviewed alongside the AI inventory.

Decision 06 · What is our reporting cadence to the board?

The board cannot govern AI from a one-off briefing. It needs a reporting cadence with the right level of summary and the right escalation triggers. In my engagements, this almost always becomes a standing Audit Committee item with quarterly KPIs and an annual deep-dive.

What good looks like

- Standing AI item on Audit Committee agenda (quarterly).
- Five KPIs minimum: inventory growth, high-risk use cases under review, supplier AI concentration, AI incidents, training completion.
- Annual narrative on AI risk appetite trajectory and forward decisions for the board.

Decision 07 · Where does the buck stop on AI incidents?

An AI incident — a model behaving in an unintended way, a supplier breach involving AI features, an automated decision causing customer harm — has a different escalation path than a traditional cyber incident. The board needs to make sure that path exists, that it has a named accountable executive, and that it tests at least once a year.

What good looks like

- AI incident playbook integrated with the wider IR plan; named accountable executive (usually the AI owner from Decision 01).
- Notification thresholds defined (regulators, customers, board, insurer).
- Annual tabletop including at least one AI scenario.

03 · WHERE ARE YOU NOW?

The AI governance maturity model.

Most UK boards I assess in 2026 are at level 1 (Ad hoc) or level 2 (Reactive). Level 3 is the threshold of "defensible". The EU AI Act, ISO 42001 and emerging customer assurance expectations are pulling the bar up to level 4.

The AI Governance Maturity Model

Five levels. Most UK boards I assess are at level 1 or 2.

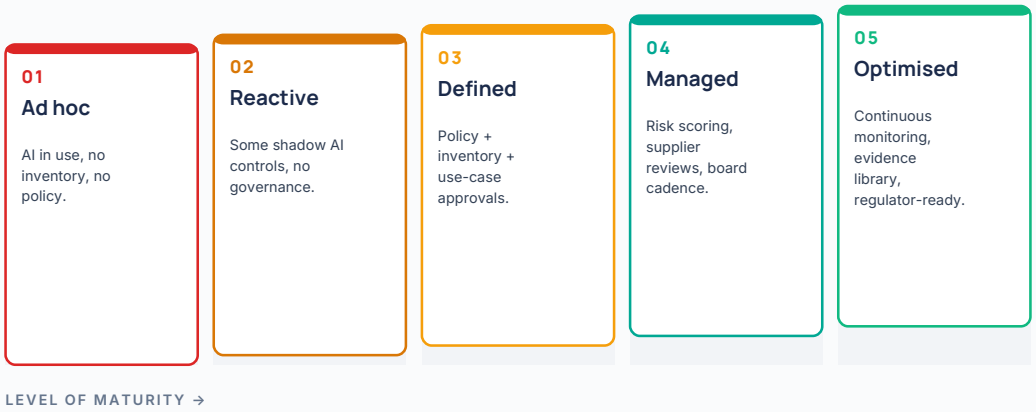

The honest position: most boards are still at the start. The good news is the path is short and well-mapped.

> **Self-assessment.** Pick the level that best describes your organisation today, then identify the single biggest gap blocking you from the next level. That gap is your next 90-day priority.

The use-case approval decision tree.

Once the seven decisions are made, the operating cadence runs on a single decision artefact: a use-case approval flow. New AI use cases come into the inventory, are screened against four checks, and end up in one of four buckets: approve, monitor, conditional approve, or block.

Use-case approval decision tree

Five decision points. Each branch ends in approve, conditional approve, monitor, or block.

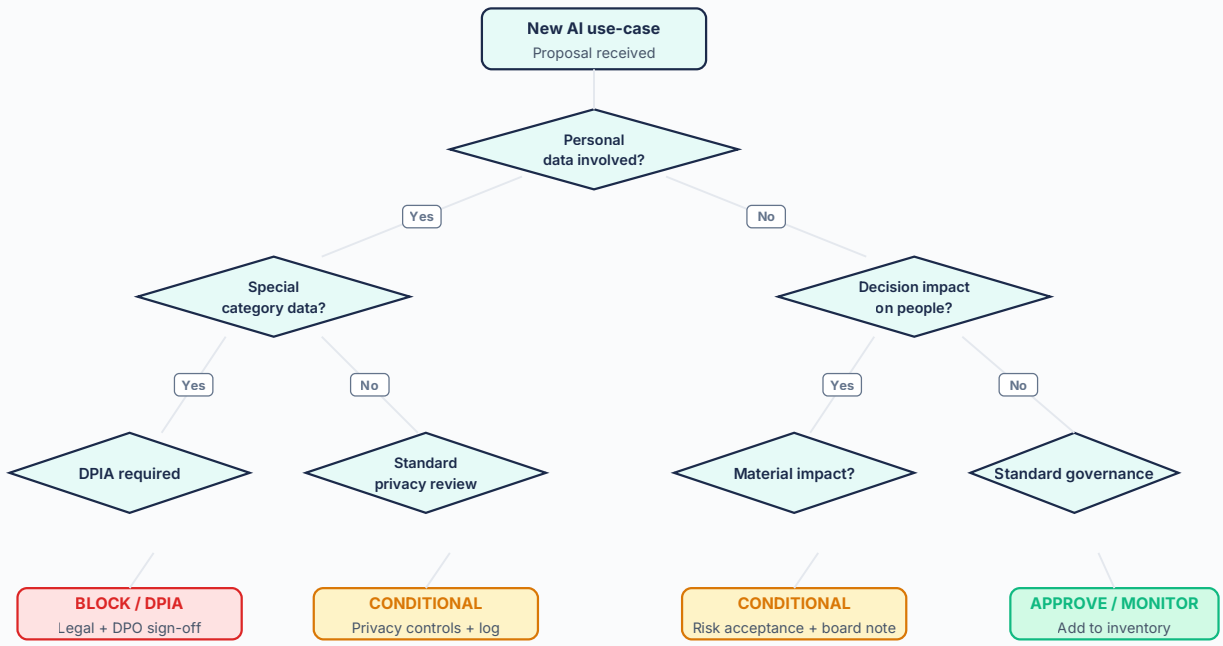

The decision flow most UK organisations should use as their default. Edge cases are escalated; the routine 80% are approved or conditioned in minutes.

05 · WHAT I SEE GO WRONG

The five most common board-level AI governance mistakes.

From advisory engagements across financial services, telecommunications, energy and the public sector, these are the patterns that come up most often. None of them are esoteric. All of them are correctable in 90 days.

THE MISTAKE	WHAT IT LOOKS LIKE	HOW TO FIX IT
Ownership lives in the CISO's inbox	The CIO and the DPO assume the CISO will absorb AI governance. The CISO has neither the mandate nor the resourcing.	Board names a single accountable executive in the minutes; RACI signed by Audit Committee.
The inventory is a slide, not a register	"We have an AI inventory" turns out to be a one-off slide deck from a 2024 risk workshop.	One canonical register, refreshed quarterly, owned by Decision 01's executive.
Risk appetite collapses AI into "technology risk"	The 2023 risk appetite statement does not mention AI. Decisions are made ad hoc.	One-page AI-specific appetite statement signed by board, with explicit positions.
Shadow AI is "addressed by the policy"	An acceptable-use policy exists; no DLP / network / monitoring controls back it; no usage survey.	Pair policy with at least one technical control and at least one usage survey per year.
Supplier AI risk is treated as "AI vendors" only	Microsoft, Google, Salesforce — major ICT providers — quietly adding AI features go un-reviewed.	Apply AI risk lens to all material ICT suppliers; tag and review at contract renewal.

If you would like senior support across any of this.

InfoSecAI is an independent UK consultancy providing fractional CISO, vCISO, AI governance and security advisory services. We help boards, executives and security leaders deliver practical governance, controls and operating-model change across information security, GRC and AI.

20+
years senior security experience

7
sectors regularly served

£12m
largest cyber programme delivered

Operationalise the seven decisions in 90 minutes.

The InfoSecAI AI Governance & Privacy Readiness Copilot turns this brief into a guided web workspace.

- **Ten guided modules** — use-case capture, shadow AI, privacy, security, supplier review, governance, scoring, action plan, board summary, evidence.
- **AI-assisted drafting** of board summaries, action plans and supplier review notes from your structured inputs.
- **Brand-locked exports** ready for direct audit committee circulation.

Live at **infosecai-aig-copilot.vercel.app**.

To talk about your AI governance position

A 30-minute conversation about where your organisation sits on the maturity model, the one or two decisions most worth making first, and whether a fractional or interim CISO model is the right shape. No charge, no obligation.

Email paul.jolliffe@infosecai.net or book directly: infosecai.net.

Founder credentials: MBA (Henley Business School) · CISSP · ISO 27001:2022 LA / LI / IA · PRINCE2 Practitioner. Past engagements include IBM, KPMG, PwC, T-Systems (Deutsche Telekom), Philip Morris International, Britannia Financial Group, MTN, Phoenix Software.