InfoSecAI

A BRIEF FOR CISOS · BISOS · GRC LEADERSHIP

WHITE PAPER · PAPER 5 · 2026

The Board Pack for AI Assurance

A practitioner's brief for board chairs, audit and risk committee chairs, CEOs, CISOs, general counsel and the executives who prepare the board pack on artificial intelligence.

AUTHORED BY

Paul Jolliffe

By Paul Jolliffe, Founder and Director, InfoSecAI Limited · MBA · CISSP · ISO 27001 Lead Auditor

01 · EXECUTIVE SUMMARY

Executive summary

The artificial intelligence (AI) question facing boards has moved. It is no longer "are we using AI". It is "can we trust the AI we are using, and can we prove it". The shift sounds incremental. In practice it changes what the executive team has to put in front of the board, in what form, and on what cadence.

This brief is the fifth and final paper in InfoSecAI's executive series, From AI Ambition to AI Assurance. The first four covered the operating model, the discovery of the AI estate, the control plane for agentic AI, and the transition from pilot to scale. This paper sets out what the board should see, how often, and what to do with it.

The thesis is direct. The board does not need a forty-slide AI strategy update. It needs a clear answer to four questions: what AI are we using, what could go wrong, who owns it, and what evidence proves it is controlled. Anything that does not answer one of those four questions belongs in an executive committee paper, not on the board agenda.

The distinguishing phrase across the series stands as the close: AI assurance evidence, not AI reassurance narrative.

02 · WHY THIS MATTERS NOW

Why this matters now

Three forces are pushing AI from the strategy agenda onto the assurance agenda.

- **Regulatory.** The European Union Artificial Intelligence Act (EU AI Act, Regulation (EU) 2024/1689) applies the bulk of its high-risk system obligations from 2 August 2026. Article 4 has imposed AI literacy obligations on providers and deployers since 2 February 2025. The Information Commissioner's Office (ICO) in the United Kingdom signals the same accountability expectation. Board oversight that cannot evidence AI literacy and active risk decisions is now a supervisory exposure.
- **Commercial.** Major customers and procurement functions are asking AI-specific due-diligence questions. Suppliers are being asked for AI inventories, model cards, training-data representations, deployer documentation and incident history. Boards that cannot demonstrate they oversee these artefacts find their organisations losing pipeline.
- **Operational.** The pattern across the four prior papers in this series is consistent: AI risk is real, AI value is uneven, and the gap between adoption and assurance is widening. Boards that receive AI strategy slides without AI assurance evidence are governing on belief, not on data.

The combined position is that AI now belongs on the board pack in a form the board can act on. Not as a thirty-page innovation update; as a one-page assurance dashboard with the underlying evidence available on request.

03 · THE FOUR BOARD QUESTIONS

The four board questions

Four questions structure the board's AI oversight. Every paper, every dashboard, every committee discussion answers one of them or it does not belong.

- **What AI are we using?** Current inventory of AI systems in scope: in-house models, AI features inside procured software, agents in evaluation, agents in production. Counted, classified by EU AI Act risk tier, tagged by data sensitivity, refreshed on a documented cadence.
- **What could go wrong?** The named risks the organisation has accepted, mitigated or escalated. Top five with status, control owner and most recent evidence. Risks that materialised in the period, classified as incidents or near-misses.
- **Who owns it?** The accountable executives for the AI estate, the operating model, supplier relationships and response capability. Named individuals, not committees.
- **What evidence proves it is controlled?** Operational artefacts, not policy documents. Logs, control evidence, assurance reports, supplier attestations, incident records, decision logs, scorecard refreshes. Available on the board portal with a sample reviewed once per cycle.

Boards that ask the four questions consistently shift the executive conversation in three quarters. Boards that ask "what is our AI strategy" do not.

04 · THE BOARD AI ASSURANCE DASHBOARD

The board AI assurance dashboard

The dashboard is the one-page artefact that anchors the AI item on every standing board agenda. It is the page the board reads first, and the page that decides whether the rest of the pack is needed. Eight panels, sized so the page is readable at a glance.

Three design rules govern the dashboard. Each panel carries a single number or a short list, never narrative. The narrative sits in the supporting pack, indexed against the panel that triggered it. Each panel has an explicit evidence reference; if the reference cannot be produced, the panel cannot be coloured green. The panel labels are stable from quarter to quarter; only the data changes, so the board can read trends without re-orienting.

The example dashboard below is the format InfoSecAI uses with mid-sized and large organisations. The numbers are illustrative. The layout, the panel selection and the four-question alignment at the foot are the operating discipline.

FIGURE 1 · The one-page board AI assurance dashboard.

Eight panels. Quarterly board cadence, monthly executive committee refresh. Each panel answers a question.

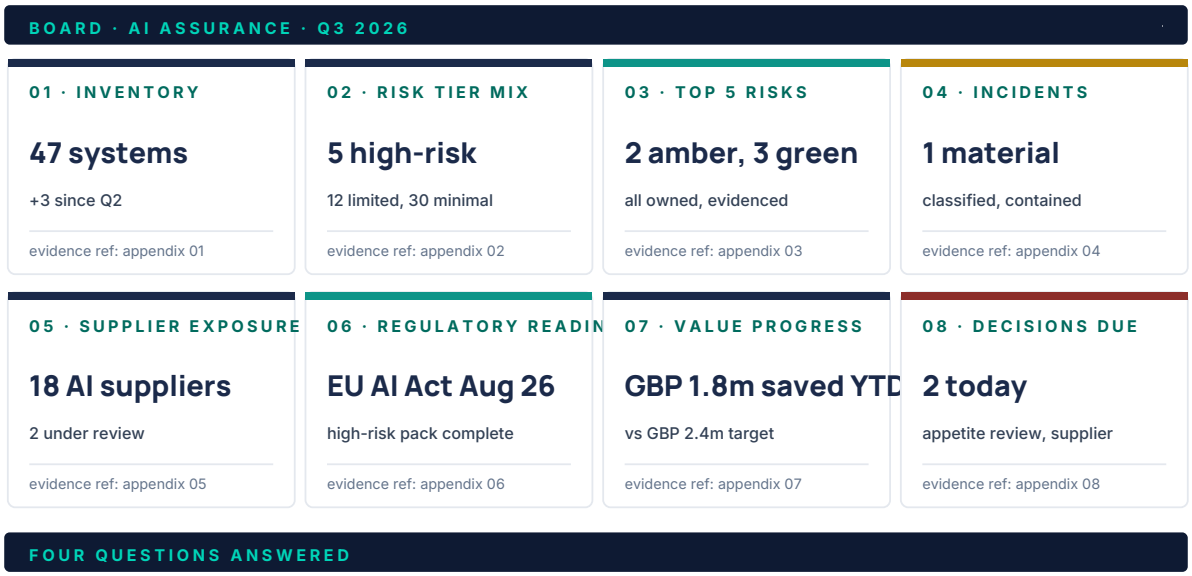

Each panel carries a single number or short list, never narrative. The narrative sits in the supporting pack, indexed to the panel.

- **Inventory.** Total AI systems in scope, with movement against the prior period.
- **Risk tier mix.** Distribution across EU AI Act tiers, including any high-risk systems with count and trend.
- **Top risks.** The five highest current AI risks with owner, status and most recent evidence reference.
- **Incidents.** Material AI incidents in the period, with classification under the operating model and the response status.
- **Supplier exposure.** AI suppliers with material exposure, count under review, any with documented deficiencies.
- **Regulatory readiness.** Position against the applicable regulatory framework, with the next milestone date.
- **Value progress.** Portfolio value position against agreed targets, drawn from the transformation portfolio in Paper 4.
- **Decisions due.** The named decisions the board is being asked to make, with owner and required-by date.

The dashboard is the executive's commitment that the four questions are answered. Failure to populate a panel is itself a signal.

05 · THE RISK APPETITE SPECTRUM

The risk appetite spectrum

The board owns the risk appetite. The executive operates within it. For AI, the appetite is best expressed as a spectrum across the four operating modes the agent paper in this series defined: observe, advise, act with approval, act autonomously.

FIGURE 2 · AI risk appetite spectrum.

Four operating modes. Each tagged with the default board position. Reviewed annually.

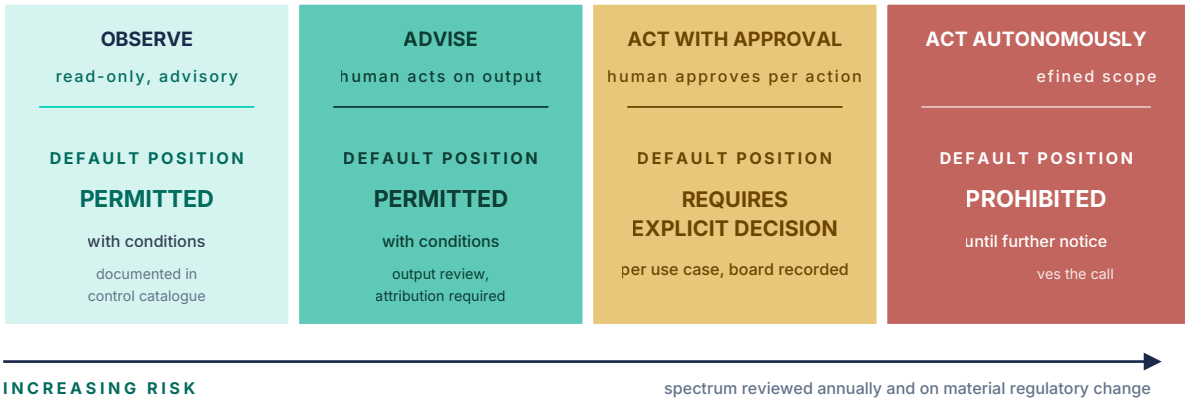

Most boards land at "permitted with conditions" for observe and advise, "requires explicit decision" for act with approval, and "prohibited until further notice" for act autonomously. That is a defensible position for 2026. The spectrum is revisited annually and on material regulatory or supervisory change.

The discipline is to write the appetite down, not infer it. An organisation whose appetite lives in conversation rather than on paper cannot demonstrate to a regulator, an auditor or a customer that the executive is operating within it.

06 · THE AI ASSURANCE EVIDENCE PYRAMID

The AI assurance evidence pyramid

Assurance is not a single artefact. It is a pyramid that the board can interrogate to whatever depth a question requires. The board does not read every layer at every meeting. It reads the dashboard, then drops to the layer that answers the question being asked.

The pyramid is not theoretical. When a board asks "are we ready for the EU AI Act August 2026 milestone", the answer comes from the policy layer (we have a position), the control layer (we have controls mapped to the obligations), the evidence layer (the controls are operating with logs to show it), and the assurance layer (internal audit has tested the controls, or our ISO/IEC 42001 attestation covers them). A pack that can answer only at the top of the pyramid is not yet a board pack.

FIGURE 3 · The AI assurance evidence pyramid.

Four layers. The board can interrogate to any depth. The bottom carries the weight.

BOARD-PACK DISCIPLINE

A pack that draws from all four layers is materially different from one that draws only from the top.

- **Policy.** The set of statements about what is and is not permitted. Necessary, but the smallest layer.
- **Control catalogue.** The defined controls that operationalise the policy: access, logging, supplier review, intake, kill-switch, incident response. Owned, mapped to the AI estate, refreshed quarterly.
- **Operational evidence.** The artefacts that prove the controls are operating. Logs, reviews, decision records, scorecards, incident reports. The bulk of the assurance load sits here, and this is where boards that ask the right questions go looking when the dashboard surfaces a signal.
- **Independent assurance.** Internal audit reports, external attestation under ISO/IEC 42001:2023, supplier attestations, regulatory reviews. The smallest layer in volume, the heaviest in weight when the board is asked to defend a decision.

A board pack that draws from all four layers is materially different from one that draws only from the top.

07 · THE REPORTING CADENCE

The reporting cadence

Board oversight of AI has three reporting paths. They run at different speeds and they answer different questions. A board operating only at the slowest speed is governing on a quarterly snapshot of an estate that changes weekly. A board operating only at the fastest speed is receiving operational management work that should have been resolved by the executive.

Each path has its own trigger, owner, content shape and decision authority. The CISO is accountable for the integrity of the cadence. The company secretary protects the board cadence. The chair of the audit or risk committee owns the on-incident path. Without

these named roles, the paths blur into each other and the board ends up either under-informed or operationally burdened.

FIGURE 4 · The board AI reporting cadence map.

Three paths. Different speeds. Different decision authority. All three are needed.

PATH	TRIGGER	OWNER	CONTENT	DECISION AUTHORITY
QUARTERLY BOARD 1 of 3	standing agenda	company secretary	dashboard + pack	appetite, scope, prohibited
MONTHLY EXEC CTTE 2 of 3	rolling cycle	CISO + COO	dashboard, decision queue	operational change
ON-INCIDENT 3 of 3	material incident	CISO + audit chair	incident brief, response	pause, contain, notify

PATH DISCIPLINE

Quarterly only is too slow. Monthly only pushes operational work to the board. Three paths scales.

- **Quarterly board.** The standing AI item on the board agenda. Dashboard plus the supporting pack. Decisions taken at this cadence include risk appetite, prohibited uses, investment thresholds and assurance scope.
- **Monthly executive committee.** The operating cadence. Dashboard refresh, portfolio review, decision queue for items below the board threshold. Where most material change is authorised.
- **On-incident escalation.** Triggered by material AI incidents under the operating-model classification or by external events that change the regulatory or supplier position. Path runs through the chief information security officer (CISO) and the chair of the audit or risk committee, with the board notified by the next standing meeting or earlier where the matter requires.

Boards that try to operate AI through the quarterly cadence alone are slower than the change rate of the estate. Boards that try to operate through monthly escalation are receiving operational management work and will eventually push back. The three-path cadence is the structure that scales.

08 · INFORMATION SECURITY IMPLICATIONS

Information security implications

Six integration points connect the board pack to the security control plane. The chief information security officer (CISO) is the executive who maintains them.

- **Inventory accuracy.** The AI system register feeding the dashboard is the same register the operating model produces. Discrepancies between what the board sees and what the security team operates indicate a control gap.

- **Incident classification.** AI incidents are classified using the operating-model standard operating procedure introduced in Paper 1, including the agentic-AI branch from Paper 3 and the shadow AI branch from Paper 2. Incidents that bypass classification create a board-level blind spot.
- **Supplier disclosure.** Supplier exposure draws from contractual representations, attestation evidence and breach history. The supplier review process is the source of truth.
- **Logging and reconstructibility.** Every dashboard signal must be traceable to operational evidence the security team can produce on request. If the security team cannot reconstruct what the AI did, the dashboard is a story.
- **Privacy and data protection alignment.** The data protection officer signs off the privacy panel and the data exposure summary, in the language the Information Commissioner's Office uses, not in marketing language.
- **Resilience.** The continuity panel reflects the manual fallback for AI-enabled processes. AI-enabled processes without a documented fallback are operational risks the board pack must surface.

The integration is the test. A dashboard that lights up green when the underlying controls are amber is more dangerous than no dashboard at all.

09 · EXECUTIVE DECISIONS AND THE 30-DAY BOARD-READINESS SPRINT

Executive decisions and the 30-day board-readiness sprint

The six decisions the board must take

Six decisions reliably move oversight from update to assurance. They do not get delegated.

- **Who is the accountable executive for AI risk?** One named individual, typically the CISO, with a defined delegate for operations.
- **What is the risk appetite spectrum?** Written, dated, signed by the chair, reviewed annually.
- **What goes on the standing agenda?** A standing AI item on the board agenda with the dashboard as its anchor. Recorded in the terms of reference.
- **What is the escalation threshold for incidents?** A bright-line classification, with the CISO authorised to escalate without waiting for the next meeting.
- **What is the AI literacy expectation for board members?** Defined and recorded under Article 4 of the EU AI Act, with a documented programme to meet it.
- **What is the evidence pack standard?** The pyramid layers expected behind each dashboard panel, with a sample reviewed once a year.

The 30-day board-readiness sprint

The sprint that converts a board pack from strategy slides to assurance evidence runs in four ordered weeks, shown in Figure 5. The structure does not require new platforms or new vendors; it requires five named people, a shared standard for the dashboard, and the chair's attendance at the dry run.

FIGURE 5 · The 30-day board-readiness sprint.

Four ordered weeks. Each week has an owner and an output. The board reads the dashboard at the end of week 4.

WEEK	ACTIVITY	OWNER	OUTPUT
W1 ALIGN	agree the four questions, dashboard structure, cadence	chair, CEO, CISO, co-sec	board oversight statement
W2 POPULATE	build dashboard v1 against the four questions	CISO + exec team	dashboard v1 + gap list
W3 DRY RUN	run dashboard through ExCo with chair attending	CEO and chair	gap-closed dashboard
W4 EMBED	add to standing board agenda, update ToR	company secretary	first quarterly cycle live

SPRINT DISCIPLINE
No new platforms. No new vendors. Five named people, four ordered weeks, one decision per week.

Worked example: Marchwood Health Trust

Marchwood Health Trust, a mid-sized National Health Service mental-health trust, ran the sprint in May 2026 ahead of its July board. The trust had nine AI deployments live across clinical triage support, administrative automation, fraud detection and supplier-embedded features in the electronic patient record.

The first dashboard surfaced three gaps. Two clinical triage deployments lacked documented human oversight design. One administrative deployment was unregistered. The supplier exposure panel could not be populated because procurement did not have a complete view of which contracts included AI features.

The board took three decisions in July. The two clinical deployments were paused pending an oversight design under EU AI Act Article 14. The unregistered deployment was registered and controlled within the month. Procurement was given four weeks to produce a complete AI supplier view, chair-sponsored.

By October, all nine deployments were registered, the supplier view was complete, the clinical deployments had resumed under a documented oversight model, and AI moved from an annual deep-dive to a standing board item.

10 · QUESTIONS EVERY BOARD MEMBER SHOULD BE READY TO ASK

Questions every board member should be ready to ask

- For the three highest-risk AI systems on the dashboard, what is the named owner, the most recent evidence of control operation, and the supplier attestation position?
- Which decision is being asked of the board today, and what is the recommended decision with the reasoning?
- What changed in the AI estate in the period covered, and what is the trend over the last four periods?
- For any incident in the period, what was classified, what was reported externally, and what was the response time against the standard operating procedure?
- What is the single gap on the dashboard that has been pending the longest, and what evidence would close it? That gap is the priority for the next executive committee.

11 · CLOSING THOUGHT

Closing thought

The five papers in this series have laid out one argument across five surfaces. The operating model is the foundation. Discovery of the estate is the front end. The control plane is the runtime. Pilot-to-production discipline is the value engine. The board pack is the oversight layer that proves the rest is operating.

The thread that connects them is one phrase: AI assurance evidence, not AI reassurance narrative. A board that asks the four questions and an executive that answers with operational evidence is what trusted enterprise AI looks like.

AI adoption is no longer the hard part. Governing, securing, evidencing and scaling it is. That is the close of the series. Thank you for reading.

12 · SOURCE REGISTER

Source register

All sources verified to primary publisher on 5 June 2026.

#	SOURCE	USE	LINK
1	EU AI Act (Regulation 2024/1689)	Article 4 literacy, Article 14 oversight, Article 26 deployer, Article 73 incident reporting	https://eur-lex.europa.eu/eli/reg/2024/1689
2	ISO/IEC 42001:2023 AI Management System	Independent assurance layer of the evidence pyramid	https://www.iso.org/standard/42001

#	SOURCE	USE	LINK
3	NIST AI RMF 1.0	Govern, Map, Measure, Manage in board oversight	https://www.nist.gov/itl/ai-risk-management-framework
4	NIST AI 600-1 Generative AI Profile (2024)	GenAI specific risks for the board risk register	https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
5	NCSC Secure AI System Development Guidelines	Secure operations framing for the dashboard	https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development
6	ICO AI guidance and AI audit framework	UK supervisory expectation on board oversight	https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/
7	European Commission AI Act regulatory framework	Provider/deployer determination at board level	https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

13 · ABOUT THIS SERIES

About this series

From AI Ambition to AI Assurance is a five-paper executive briefing series, 1 to 5 June 2026.

1. AI Governance Is No Longer a Policy Problem
2. The Shadow AI Exposure Map
3. Securing Agentic AI Before It Acts
4. Why AI Transformation Fails After the Pilot
5. The Board Pack for AI Assurance (this paper)

Each paper is published as a 12-page executive briefing under the InfoSecAI Blog Template. The full series is available at infosecai.net/insights for subscribers to the InfoSecAI insights list.

14 · PRACTITIONER NOTE

Practitioner note

This briefing is practitioner interpretation, not legal advice. For regulated deployments, validate final claims against current legal obligations, sector-specific requirements and the original primary sources before relying on them.

InfoSecAI

About InfoSecAI

InfoSecAI is an independent UK consultancy helping organisations turn security, regulatory, resilience and AI governance requirements into practical operating models, stronger controls and robust delivery.

We work across strategy, governance, risk, compliance, AI security, assurance, operations and engineering. Our services help leadership teams assess their current position, align to standards and regulation, define the target operating model, and deliver the governance, controls, artefacts and ways of working needed to move from intent to implementation.

Our toolkit capability accelerates structured work across ISO 27001, ISO 22301, ISO 42001, NIST CSF, NIST AI RMF, CIS Controls, Cyber Essentials, DORA, NIS 2, the EU AI Act, GDPR, UK GDPR, SOC 1 and SOC 2. The approach combines AI-enabled workflow support with senior practitioner judgement, so outputs remain proportionate, usable and connected to the way the organisation actually operates.

InfoSecAI was founded in **2025** by **Paul Jolliffe**. The company is built for organisations that need clarity, senior leadership and hands-on delivery across information security and AI governance, without adding unnecessary complexity or treating compliance as a paperwork exercise.

infosecai.net · paul.jolliffe@infosecai.net