InfoSecAI

A BRIEF FOR CISOS · BISOS · GRC LEADERSHIP

WHITE PAPER · 2026

The AI Governance Gap

A board-facing brief on why AI adoption has outpaced AI governance, and the minimum credible posture for 2026.

AUTHORED BY

Paul Jolliffe

By Paul Jolliffe, Founder & Director, InfoSecAI Limited · CISSP · ISO/IEC 27001:2022 Lead Auditor · MBA · 20+ years across financial services, telecommunications, energy, healthcare, technology, and public sector

01 · EXECUTIVE SUMMARY

Executive summary

Most organisations now have artificial intelligence (AI) running ahead of AI governance, and the distance between the two is widening. Employees are using third-party generative AI to draft customer communications, summarise contracts, write code, and reason over sensitive data. Vendors are quietly retrofitting AI into products the organisation already pays for. Departments are buying their own AI tools on corporate credit cards. None of this is unusual. What is unusual, and what should concern any board reading this in 2026, is that very few organisations can tell their auditor, their regulator, or their insurer what AI systems they actually have, who is accountable for them, and what controls apply.

This is the AI governance gap. It is not a single failure of policy or a missing committee. It is the structural lag between the speed of AI adoption and the speed at which most organisations bring new technology under governance. Closing it is not optional. The Information Commissioner's Office (ICO), the Financial Conduct Authority (FCA), the Prudential Regulation Authority (PRA), and the European Union (EU) supervisory authorities for the EU AI Act are converging on the same expectation, namely that the organisation can demonstrate, with evidence, that it knows what AI it is operating, what risks those systems carry, and what governance is in place to manage them.

This brief sets out the shape of the gap, why it has opened, what a minimum credible AI governance posture looks like for 2026, and ten questions the board should be putting to the executive at its next meeting.

02 · THE SHAPE OF THE GAP

The shape of the gap

The gap has three measurable features.

The first is **inventory blindness**. Ask any chief executive how many AI systems the organisation operates and the candid response is that the work is underway. The reality, on every engagement I have led in the past eighteen months, is that the first-pass inventory finds three to five times as many AI systems as the executive expected, the majority of which arrived through software-as-a-service vendor updates rather than deliberate procurement decisions. If you do not know what you have, you cannot govern it.

The second is **accountability vacuum**. AI cuts across functions in a way that few other technologies do. It touches data, security, privacy, customer experience, human resources, finance, and product. In most organisations no single executive owns it end-to-end. The Chief Information Security Officer (CISO) owns the security risk, the Chief Risk

Officer (CRO) owns the regulatory risk, the Data Protection Officer (DPO) owns the privacy risk, the Chief Technology Officer (CTO) owns the engineering risk, and nobody owns the aggregate. The board sees the parts and never sees the whole.

The third is **control debt**. The controls the organisation already runs, namely access management, change management, third-party risk, model validation, data classification, and incident response, were designed for systems that behave deterministically. AI systems do not. They produce different outputs for the same input, they degrade over time, they can be manipulated by adversarial inputs, and they leak training data in ways that traditional data loss prevention does not detect. The existing control stack still applies, but it is not sufficient on its own.

03 · WHY THE GAP HAS OPENED

Why the gap has opened

The gap is not a story about negligence. It is a story about speed, structure, and standards arriving in the wrong order.

The speed is obvious. The generative AI wave that began in late 2022 reached enterprise adoption faster than any previous technology cycle. Procurement processes designed for eighteen-month evaluation periods cannot keep up with a vendor refresh cycle measured in weeks.

The structure is less obvious. The existing governance architecture in most organisations was built around three lines of defence (3LoD) that assume the first line can describe the asset, the second line can evaluate the risk, and the third line can audit the controls. AI breaks the first assumption. A foundation model running inside a vendor product is not an asset the first line can fully describe, because the model itself is opaque and the vendor's documentation does not surface what the second line needs to know.

The standards arrived after the wave. ISO/IEC 42001:2023 was published in December 2023. The EU AI Act entered into force in August 2024 with staggered application dates running through 2027. The National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF) version 1.0 was published in January 2023, and the Generative AI Profile (NIST AI 600-1) followed in July 2024. The UK's regulatory approach is still pluralistic, with the ICO, the FCA, the PRA, the Competition and Markets Authority (CMA), and Ofcom each applying their existing remits. An executive cannot reasonably be expected to have built a programme against a standards landscape that did not yet exist.

InfoSecAI

A BRIEF FOR BOARDS AND AUDIT COMMITTEES

The AI Governance Gap

Why AI adoption has outpaced AI governance, and the minimum credible posture for 2026.

3 to 5×	Aug 2026	6	90 days
First-pass AI inventory undercount on every engagement	EU AI Act high risk obligations begin to apply	Structural moves to a defensible governance posture	From this board meeting to a defensible position

THE SHAPE OF THE GAP

Three measurable features.

The gap is structural, concrete, and addressable, but only once you decide to look.

Inventory blindness

FEATURE 01

First-pass inventories find three to five times more AI systems than expected, mostly from vendor software updates.

Accountability vacuum

FEATURE 02

Four executives own parts of AI risk. Nobody owns the aggregate.

Control debt

FEATURE 03

Existing controls were built for systems that behave deterministically. AI does not.

REGULATORY LANDSCAPE

Regulatory application timeline, 2022 to 2027.

CLOSING THE GAP

Six moves. Ninety days. No transformation programme.

01 Live AI inventory
Every system, including vendor embedded AI. Refreshed quarterly.
ARTEFACT Canonical AI register

02 Named senior executive
One executive owns the aggregate and reports to the board.
ARTEFACT Board-minuted RACI

03 Short, signed AI policy
Two to four pages, written for the workforce. Acceptable use, triggers.
ARTEFACT AI policy v1.0

04 Impact assessment
Triage every use case against the EU AI Act risk tiers.
ARTEFACT Assessment process + template

05 Risk machinery integration
AI as a dimension on existing registers. Not a new register.
ARTEFACT Consolidated risk view

06 Board cadence
AI on every board agenda until stable, then quarterly.
ARTEFACT Standing agenda + dashboard

FOR THE AUDIT COMMITTEE

Ten questions to surface the gap in one meeting.

The right answers are not yes or no. They are evidenced. If the executive cannot answer six in evidenced form, the gap is the board's problem.

1. How many AI systems do we operate, and when was that number last verified independently?
2. Who is the single named senior executive accountable for AI?
3. What is our AI policy, and how do we know employees follow it?
4. Which of our use cases would be high risk under the EU AI Act?
5. Which use cases process special category personal data?
6. When did we last conduct an AI impact assessment, and on what?
7. Which material vendors use AI inside the services they provide?
8. What is our response if a customer facing AI produces harm tomorrow?
9. What evidence would we present to the ICO, FCA or PRA tomorrow?
10. What is our risk appetite for AI, and where is it documented?

The gap is real. It is also closable.

A board that begins this work in May 2026 can be on the right side of the gap by year end.
A board that defers it will spend 2027 explaining its delay.

FULL BRIEFING AND BOARD PACK

infosecai.net

04 · WHAT THE GAP LOOKS LIKE INSIDE THE ORGANISATION

What the gap looks like inside the organisation

A worked example brings the gap into focus. Consider Brackenridge Outsourcing plc, a fictional mid-cap business process outsourcer with three thousand employees and contracts across financial services and the public sector.

In May 2026 the executive team is asked by the audit committee chair to confirm the organisation's AI exposure. The answer takes six weeks to assemble. It surfaces forty-one AI use cases. Eleven of those are formally approved, with documented business cases. Thirty are not. Of the thirty, eighteen are vendor-embedded features in tools the organisation has been using for years, namely the customer relationship management platform, the human resources information system, and the call quality assurance product. Six are generative AI assistants enabled by employees on free or personal accounts. Six are pilot projects launched by individual business units. Two are AI features the organisation paid for as part of a procurement bundle without understanding what they did.

None of the forty-one has an AI impact assessment. Three have data protection impact assessments. The CISO has visibility of nine. The DPO has visibility of seven. There is partial overlap between the two lists. The audit committee chair asks whether any of the use cases fall within the high-risk category under the EU AI Act. Nobody can answer in the meeting.

This is not a story about an unusually poorly run organisation. It is a typical 2026 inventory finding. The point is that the gap is concrete, measurable, and addressable, but only once you decide to look.

05 · THE MINIMUM CREDIBLE GOVERNANCE POSTURE FOR 2026

The minimum credible governance posture for 2026

A board does not need to commission a multi-year transformation programme to close the gap. It does need to instruct the executive to deliver a small number of specific things, in a defensible sequence, and to evidence them.

- **A live AI inventory.** Every AI system the organisation operates, including vendor-embedded features, with owner, purpose, data classification, risk tier, and review date. Refreshed quarterly. This is the single most important artefact and most organisations do not yet have one.

- **A single named senior accountable executive for AI.** Whether that is the CISO, the CRO, the Chief Data Officer (CDO), or a newly appointed Chief AI Officer (CAIO) is less important than the fact that one person owns the aggregate and reports to the board. Federated ownership has had three years to work and has not.
- **An AI policy that is short, signed, and known.** Two to four pages, written for the workforce rather than the lawyer, addressing acceptable use of public generative AI, the procurement gate for new AI tools, mandatory impact assessment triggers, and the escalation route for AI incidents. Aligned to ISO/IEC 42001:2023 Clause 5.2, the NIST AI RMF Govern function, and the EU AI Act transparency obligations.
- **An AI impact assessment process.** A proportionate triage that classifies every new and existing AI use case as prohibited, high-risk, limited-risk, or minimal-risk, and routes each to the right depth of review. The EU AI Act's risk tiers are a serviceable starting point for organisations both in and out of EU scope, because they map onto the obligations regulators are increasingly likely to ask about.
- **Integration with existing risk and compliance machinery.** AI risk is not a new register. It is a new dimension on the existing operational risk, technology risk, model risk, third-party risk, privacy, and information security registers. The board should see a consolidated view, not five separate AI papers.
- **Board cadence.** AI on the agenda at every board meeting until the inventory is stable, the policy is operating, and the impact assessment process is running. Quarterly thereafter. Without recurring board oversight the programme will drift, because every other item on the executive's agenda will demand attention first.

That is the floor. It is not the ceiling. A regulated firm in financial services will need to layer the FCA's expectations on AI in financial services, the PRA's Supervisory Statement SS1/23 on model risk management, and the Digital Operational Resilience Act (DORA) Articles 28 and 31 on ICT third-party risk and the oversight framework for critical ICT third-party providers where AI sits in a critical function. A public sector body will need to layer the Algorithmic Transparency Recording Standard (ATRS). A healthcare provider will need to layer the Medicines and Healthcare products Regulatory Agency (MHRA) software as a medical device expectations. The base posture, however, is the same in every sector.

06 · TEN QUESTIONS THE BOARD SHOULD BE ASKING

Ten questions the board should be asking

The following ten questions surface the gap in a single meeting. The right answers are not "yes" or "no" but evidenced.

1. How many AI systems do we operate, and when was that number last verified?
2. Who is the single named executive accountable for AI across the organisation?

3. What is our AI policy, and how do we know employees are following it?
4. Which of our AI use cases would fall within the high-risk category under the EU AI Act if it applied to us?
5. Which of our AI use cases process special category personal data under the UK General Data Protection Regulation (UK GDPR)?
6. When did we last conduct an AI impact assessment, and on what?
7. Which of our material vendors use AI inside the services they provide to us, and how do we know?
8. What is our incident response plan if a customer-facing AI system produces a harmful or discriminatory output tomorrow morning?
9. What evidence would we present to the ICO, the FCA, or the EU AI Office if asked tomorrow to demonstrate our AI governance?
10. What is our risk appetite for AI, and where is it documented?

If the executive cannot answer six of those ten in evidenced form, the gap is the board's problem, not a future agenda item.

07 · WHERE TO START THIS QUARTER

Where to start this quarter

The AI governance gap is closed by deliberate effort, not by a transformation programme. The first quarter is the inventory and the named owner. The second quarter is the policy and the impact assessment process. The third quarter is integration with the existing risk machinery and the first board-facing dashboard. The fourth quarter is the evidence pack the organisation can hand to its first AI assurance question, whether that comes from a regulator, a customer, an insurer, or an auditor.

A board that begins this work in May 2026 can be on the right side of the gap by year-end. A board that defers it will spend 2027 explaining to its regulator and its auditor why it did not.

The InfoSecAI AI Governance Board Pack is a free downloadable resource that gives boards the agenda, the dashboard, and the question set to start this quarter. The InfoSecAI AI Governance and Privacy Readiness Copilot is the practitioner workbench that surfaces the inventory, runs the impact assessments, and produces the evidence pack the executive will need to put in front of the board. Both are available at infosecai.net.

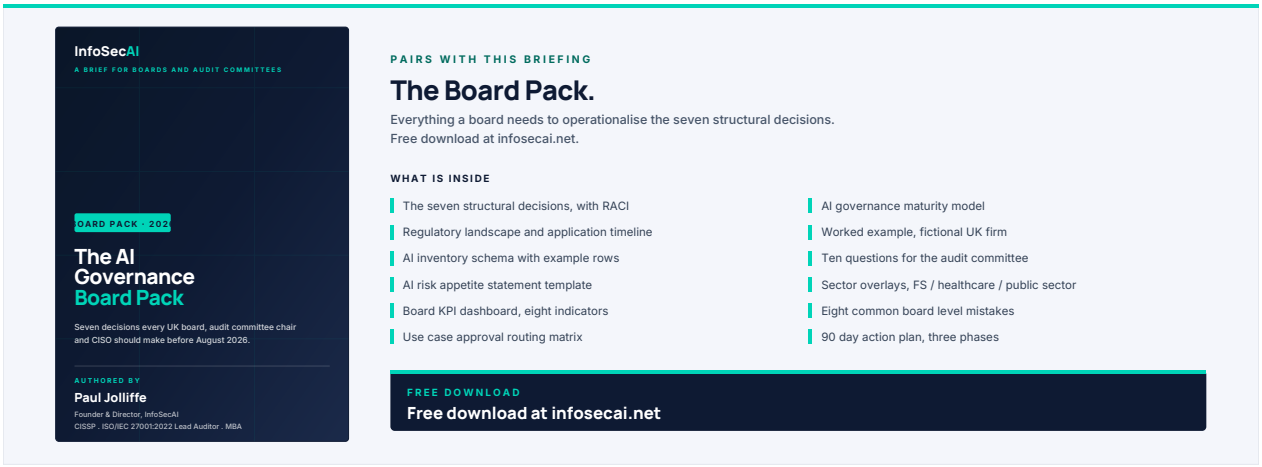

The gap is real. It is also closable. The only question is whether the board treats it as something to address this year or to be told about next year.

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