Ravinel .AI

BRIEF

AI governance in regulated industries.

What boards are getting wrong, and how to fix it before the regulator asks.



We make AI scale through people.

In regulated industries, AI is already making or shaping decisions that matter: who gets approved, what gets flagged, which case is handled first. Most boards know this in principle. Far fewer can answer three simple questions: what AI do we use, how risky is each system, and who is accountable when it goes wrong?

That gap is about to matter more. Across many markets, AI regulation is moving from voluntary principles to supervision, with expectations on inventories, risk classification, human oversight and reporting. The organisations that close the gap now will scale AI faster later. The others will spend the next two years explaining themselves.

Five things boards are getting wrong

01 **Treating governance as a policy document.**

A responsible AI policy approved once a year does not govern anything. Governance lives in the decisions made every week: what gets built, what gets bought, what goes live, and what gets stopped.

02 **Not knowing what AI is in use.**

Most organisations can list their flagship AI projects. Very few can list the AI inside vendor software, or the generative AI tools staff use every day. You cannot govern what you cannot see.

03 **Managing every system the same way.**

A chatbot that answers opening hours and a model that influences approvals do not carry the same risk. Applying the same heavy process to both slows everything down, and still misses what matters.

04 **Leaving accountability with IT.**

When AI sits only with the technology function, business owners stop feeling responsible for the outcomes it produces. Accountability must sit with the people who own the decision the AI supports.

05 **Looking only before launch.**

AI systems change after they go live. Data drifts, vendors update models, people use tools in new ways. Without monitoring, the risk assessment done at launch is out of date within months.

What good looks like: govern the lifecycle, not the document

Effective AI governance follows every system through the same five steps, and back again whenever it changes.

1

Context.

What the system does, for whom, with which data, and where it must not be used.

2

Identify.

The risks to people, to the business, to data, to security and to compliance.

3

Assess.

Likelihood and impact, and a clear risk level: unacceptable, high, limited or minimal.

4

Treat.

Controls and human oversight in proportion to the risk, and a clear decision: approve, limit or stop.

5

Monitor.

Performance, drift, fairness and incidents, with regular review.

The effort follows the risk. Minimal-risk tools move quickly with light checks. High-risk systems get deeper testing, named owners and closer oversight. Nothing unacceptable goes live.

A practical first 90 days

Days 1–30: see it.

Build one inventory of every AI system in use, including vendor and generative AI tools, each with a named business owner.

Days 31–60: sort it.

Give every system a risk level using one simple, documented method. Agree who can approve what, and set up a small AI committee with business, risk, legal and technology.

Days 61–90: run it.

Put a risk assessment and sign-off in front of every higher-risk launch. Start monitoring the systems already live, and send the board its first short AI risk report.

Ninety days will not make governance perfect. It will make it real, and that is what regulators, customers and boards look for first.

Five questions every board should ask this quarter

01	Do we have a complete inventory of the AI we use, including AI inside vendor products?
02	Which of our AI systems are high-risk, and who owns each one?
03	What has to happen before a high-risk AI system goes live, and who signs it off?
04	How do we know our AI systems still work as intended after launch?
05	When did the board last receive a report on AI risk, and what did it say?

If the answers are unclear, that is the starting point, not a failure.

Where to start

Take the free AI Governance and Risk Readiness check: twelve questions, about six minutes, with your top three next moves. Or book a 30-minute diagnostic with Ravinel.

This brief is general information, not legal advice. Aligned with ISO/IEC 42001, the NIST AI Risk Management Framework and the EU AI Act.

We make AI scale through people.

Govern the lifecycle. Not the document.

WEBSITE
www.ravinel.ai

EMAIL
info@ravinel.ai

GOVERNANCE CHECK
ravinel.ai/ai-tools/governance-risk-readiness

We make AI scale through people.