QuantSpark

How regulated businesses shift their governance frames to control AI, reduce risk, and scale with confidence

AI at Scale in a Regulated World

Compliance is the floor. Orchestration is the control layer that keeps pace with AI as it scales.

An enterprise can be fully compliant and still be fragmented. Policies describe desired behaviour, but orchestration enforces it.

13x

more likely to be scaling AI

Organisations using orchestration-led governance vs compliance-only approaches

IBM, 2024

12%

of organisations have orchestration platforms in place

The gap between policy and control is the default, not the exception

IBM, 2024

6x

greater productivity impact

Orchestration-led governance vs compliance-only organisations

IBM, 2024

Contents

Executive Summary

The case in 60 sec

Artificial intelligence is no longer a side project for regulated businesses. It is becoming part of core operations, customer service, decision support, compliance workflows, and third-party ecosystems, which means the governance model must evolve as quickly as the technology itself.

IBM's research shows that organisations with orchestration-led governance are 13 times more likely to be scaling AI, and they report 6 times greater productivity impact than organisations relying on compliance-only approaches.

For compliance-led organisations, the immediate priority is how to govern AI continuously. That requires moving from policy-led oversight to an operational control layer that connects inventory, identity, risk management, auditability, and enforcement inside the AI estate.

The distinction that matters is between organisations that can govern AI continuously (across the full lifecycle, all systems, and every third-party integration) and those that cannot. The first group is scaling with confidence.

The second is accumulating exposure without knowing it.

Technology is no longer the limiting factor, and the IBM research is unambiguous about the cost: organisations relying on compliance-only approaches are

6x less likely to see meaningful productivity impact from AI, and they carry substantially higher costs when irregularities occur.

THE CASE FOR ACTION

The audience is every General Counsel, Chief Risk Officer, Chief Compliance Officer, and Chief Operating Officer in a regulated business who has signed off on an AI policy without yet asking whether that policy is enforceable in production.

Section I

The Signal: AI Is Scaling Faster Than Governance

The signal is clear: AI is scaling faster than governance structures designed for slower, more static technologies. The visibility gap this creates is a direct exposure to regulatory, legal, and reputational risk. In compliance-heavy environments, that visibility gap is a direct exposure to regulatory, legal, and reputational risk.

The reason is simple. AI is rarely deployed as a single system. It arrives as a collection of assistants, workflows, models, and vendor tools, often adopted by different teams for different purposes, and often without a common control layer.

The underlying mechanism is important to understand because it explains why the gap does not close on its own.

Most governance frameworks for technology were designed for static deployments: a system is assessed, approved, and then monitored periodically. AI does not behave that way. Models drift, integrations multiply, and new use cases are added without formal review. Teams build workarounds when sanctioned tools feel slow. The governance surface expands continuously, and point-in-time controls lose their grip almost as soon as they are applied.

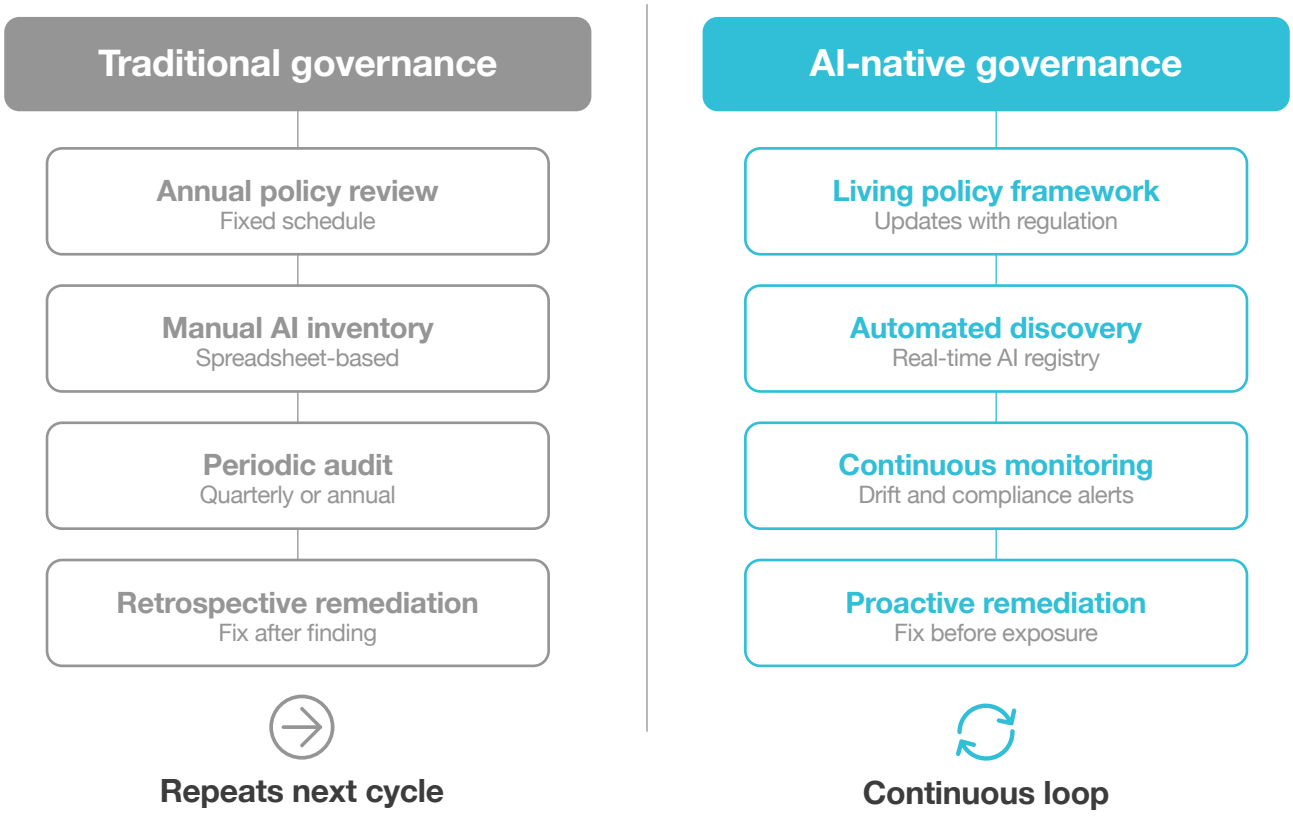

69%

of executives lack full AI visibility

IBM, 2024 — across 1,000+ senior leaders

Section II

What This Means: Compliance Alone Is No Longer Enough

Compliance is necessary but no longer sufficient

The important thing to say about compliance is that it was never designed to be the whole answer. Compliance frameworks define acceptable behaviour. They do not, on their own, enforce it in runtime. An organisation can have a fully compliant AI policy and still have models operating outside the parameters that policy describes, because the policy is a document and the models are live systems.

IBM's research draws a precise distinction: organisations with orchestration-led governance do not abandon compliance, they build on top of it. Compliance sets the standard; orchestration ensures the standard is met in practice, at the point of execution, continuously and evidentially.

Visibility is an operational imperative

When 69% of executives cannot account for the AI their organisations are running, visibility is more than just a nice-to-have. It is the prerequisite for everything else. You cannot manage what you cannot see, you cannot audit what you cannot inventory, and you cannot defend a regulatory position you cannot evidence.

IBM found that organisations with orchestration-led governance were more than twice as likely to have full visibility into their AI assets. That differential is explained by whether governance was treated as an architecture decision or as a documentation exercise.

The same visibility gap extends beyond the organisational boundary. AI arrives through vendor tools, SaaS integrations, partner systems, and outsourced processes, often without a common control layer and subject to different oversight regimes. Orchestration extends the control layer to cover this perimeter, setting enforceable standards for what third-party AI can and cannot do with data that touches the firm's regulatory exposure.

Late-stage governance creates the failures it tries to prevent

The timing of governance matters as much as its content. A compliance team can be active, well-resourced, genuinely rigorous, and still be structurally unable to prevent the problems it is reviewing, because it enters the process after the decisions that matter have already been made. Architecture choices, data pipeline design, third-party integrations: by the time these reach a governance review gate, the practical options for correction are limited and the cost of exercising them is high.

The organisations that IBM identifies as governance leaders embed controls at the point of design, making compliance an architectural property of the system rather than an assessment applied to it afterwards.

The accountability gap compounds over time

AI creates accountability challenges that are structurally different from those created by human decision-making. When a human makes a poor decision, there is a record: a meeting, an email, a sign-off. When an AI system produces an outcome that later attracts regulatory scrutiny, the ability to reconstruct the reasoning, the data inputs, the model version, and the override history depends on whether those things were logged at the time.

IBM found that orchestration-led organisations are 169% more likely to maintain transparent documentation of AI processes. The organisations that cannot demonstrate explainability in retrospect are carrying a liability that has not yet been called in.

Section III

The Research Evidence Base

The IBM data tracks outcomes at the organisational level. The regulatory and academic evidence explains the mechanism, and makes clear that the gap between governance intent and governance capability is not a niche concern. It is the central challenge of regulated AI deployment.

The IBM Evidence: Orchestration vs Compliance-Only

1,006
senior executives across

20 geographies + ***23*** industries

found a strong and consistent relationship between orchestration maturity and business performance

Organisations using orchestration-led governance were

13x

more likely to be scaling AI.

Those with a full orchestration approach saw more than

6x

the productivity impact organisations focused on compliance alone.

They also experienced **29% lower cost from AI irregularities** and **20% higher ROI on AI investments**

The research also shows that visibility and documentation are not soft benefits, they are the operational differentiators.

Organisations with orchestration-led governance were more than

2x as likely to have full visibility into their AI assets

169% more likely to maintain transparent documentation

132% more likely to protect data through anonymisation, impact assessments, and strict access controls

The FCA and Bank of England: Explainability Is a Supervisory Expectation

In April 2024, the FCA and the Bank of England jointly published updates on their strategic approach to AI. The practical implication for regulated firms is unambiguous: where AI is being used, firms should expect to need to explain that use to their regulators. Covering how risks have been identified, assessed, and managed. This is a current supervisory expectation, applied through existing frameworks including the Consumer Duty, the SM&CR, and the SYSC sourcebook requirements for governance arrangements and systems and controls.

The FCA has since launched its AI Lab and AI Live Testing initiative, with the first cohort of firms entering live testing in late 2025. Formal guidance specifically on audit trails and explainability is expected by the end of 2026. Firms that wait for that guidance to arrive before building their documentation infrastructure will be building it under scrutiny rather than ahead of it.

The PRA: Model Risk Governance Now Covers AI Explicitly

The PRA's Supervisory Statement SS1/23, effective from May 2024, established that model risk management applies across all model types, including AI and machine learning systems. It requires firms to maintain comprehensive model inventories, conduct independent validation, and implement continuous performance monitoring. The PRA has been explicit that existing MRM frameworks need to be extended, not merely reinterpreted, to cover the specific characteristics of AI: iterative development cycles, opaque decision logic, and continuous drift.

In October 2025, the PRA held dedicated CRO roundtables with 21 regulated firms specifically on AI and ML in the context of SS1/23, a clear signal that supervisory attention on this area is intensifying, not plateauing.

Gartner: The Monitoring Gap Is Structural

A Gartner survey of 360 organisations in Q2 2025 found that organisations deploying AI governance platforms are:

3.4x

more likely to achieve high effectiveness in AI governance than those that do not

Yet Gartner also projects that only

40%

of organisations deploying AI will have dedicated observability tools in place by 2028, meaning the majority are currently operating AI systems without the continuous monitoring capability that effective governance requires.

Gartner has identified this as a compounding risk: without standardised model telemetry and runtime monitoring, incident resolution for AI applications requires complex manual effort to trace and debug the behaviour of opaque models.

Section IV

The questions boards must now ask.

Strategic leadership in a regulated AI environment is about asking questions that reveal whether governance is real or performative. These are the eight questions that boards and executive committees in regulated businesses should be pressing today.

Can we produce a complete inventory of every AI system active in our business, including third-party integrations, right now?

If the answer requires more than 48 hours to compile, the inventory does not exist in any operationally meaningful sense. You cannot govern what you cannot see, and you cannot defend regulatory exposure you have not mapped.

Do our AI controls operate in real time, or do they operate at the point of review?

Policy-based governance is periodic. Orchestration-led governance is continuous. The question is whether your controls shape AI behaviour as it happens, or whether they describe the behaviour you would have preferred after the fact.

Who owns accountability when an AI system produces an outcome that attracts regulatory scrutiny?

In many organisations, AI accountability sits between functions. Nominally owned by technology, operationally driven by business lines, and reviewed by legal and compliance only when something goes wrong. That disconnect is a governance failure waiting to be triggered.

Can we reconstruct the reasoning behind any AI-assisted decision made in the last twelve months?

Explainability is not just a technical property of a model. It is a documentation discipline. Regulators in the UK and EU are increasingly treating the ability to explain AI decisions as a minimum standard.

Can you meet it today?

What proportion of AI in our business was procured through formal channels, with documented risk assessment?

Shadow AI is already present in most regulated businesses. IBM's research suggests that in many organisations, significant AI usage is unsanctioned. Every unsanctioned AI system that touches client data, financial models, or regulated processes is an unmanaged exposure.

How does our governance framework extend to the AI used by our key suppliers and outsourcing partners?

Your regulatory obligations do not stop at your firewall. If a supplier's AI system processes your clients' data or supports a regulated activity, the accountability sits with you. Do your third-party governance frameworks reflect that?

Are our legal, risk, compliance, technology, and business functions aligned on what AI can and cannot do, and is that alignment documented?

Governance by consensus is not governance. Alignment means defined decision rights, clear escalation paths, and documented approval criteria. A shared sense that we are roughly on the same page is not sufficient.

Is AI governance embedded in our delivery process, or does it arrive at the end as a review gate?

Late-stage governance is the most expensive kind. It catches problems after architecture decisions have been made, after data pipelines have been designed, and after commercial commitments have been entered into. The organisations that govern AI well build the controls in from the start.

Section V

The Regulatory Imperative

The Regulatory Framework Has Hardened

For years, AI governance in regulated industries operated in a soft-law environment: guidance, principles, expectations. The EU AI Act, the most comprehensive AI regulatory framework in the world, is now in phased enforcement.

The prohibitions on unacceptable-risk AI applications came into force in February 2025. Obligations for high-risk AI systems, including those used in credit, insurance, employment, and critical infrastructure, are active and expanding.

€35m /
7% global turnover

For organisations operating in or serving European markets, the stakes are explicit. Fines for non-compliance with the EU AI Act can reach €35 million or 7% of global annual turnover, whichever is higher. That is a boardroom-level commercial risk.

The UK Approach: Principles-Based But Not Permissive

The UK has chosen a different regulatory architecture. Sector-led, principles-based, and coordinated through the AI Safety Institute and existing regulators rather than through a single AI-specific statute. Principles-based does not mean low-accountability. It means the accountability is on the organisation to demonstrate how its AI systems deliver the outcomes regulators require.

The FCA has been explicit: AI used in regulated activities must meet the same standards of fairness, transparency, and explainability as human decisions. The PRA has identified model risk as a primary supervisory concern for AI in financial services. The ICO has issued detailed guidance on lawful AI and personal data, and is actively examining compliance.

For UK regulated businesses, whether the organisation can evidence compliance in the way that each regulator's framework demands is the non-negotiable, which increasingly means runtime evidence, not policy documentation.

The Window for Proactive Governance

There is a structural advantage available to organisations that build orchestration-led governance now, before regulatory enforcement reaches full maturity.

29%

lower cost from AI irregularities

Those organisations will arrive at supervisory review with audit trails, documented decision rights, and demonstrable controls already in place. Those that wait will be building governance infrastructure under scrutiny, at higher cost, and with less room for the iterations that operational governance always requires.

The firms that define what good looks like in this environment will shape the standards that the rest of the sector is measured against. It is the documented pattern of every regulatory maturation cycle in financial services. The early movers write the playbook, and the late movers follow it. Or fail to.

Section VI

The Execution Gap & The AiRE Response

The execution gap is where most organisations struggle. Many businesses have written policies, review boards, and approval gates, but those measures do not automatically create control at scale. IBM's research suggests that only 12% of organisations currently have orchestration platforms in place, which helps explain why AI governance remains disconnected from actual system behaviour.

In compliance-led businesses, the most common failure is late-stage governance that tries to correct design decisions after systems have already been built and deployed.

Generic AI platforms do not solve this. Microsoft Copilot and equivalent tools were not built for the domain specificity, regulatory context, and adoption inertia that define knowledge-intensive regulated businesses. The value in a law firm, a PE fund, or a financial services firm sits in the firm's own criteria, formats, and decision logic. The compliance teams in these organisations are sceptical, time-poor, and will not adopt tools that create new governance problems while claiming to solve old ones. The gap between 'we have AI tools' and 'our AI is governable' is exactly the divide AiRE was built to close.

AiRE (AI Rollout Engine) embeds alongside compliance, technology, and operations teams to design, prototype, and deploy orchestration-led AI governance inside the organisation's existing infrastructure with domain expertise, regulatory context, and adoption as the primary success metric.

The use cases below are drawn from the compliance and regulated-services environment specifically:

	AI Asset Inventory and Classification	Regulatory Obligation Mapping	Audit Trail and Explainability Logging	Third-Party AI Risk Monitoring
Current state	No single source of truth for AI systems in use; inventory compiled manually, quarterly, from departmental self-reporting. Significant shadow AI undetected.	Legal and compliance teams manually track regulatory updates across multiple frameworks (EU AI Act, FCA, ICO, PRA). Cross-referencing obligations against internal systems takes weeks per cycle.	AI-assisted decisions logged inconsistently or not at all. Reconstructing decision rationale for regulatory or legal review requires manual forensic work across multiple systems.	Supplier AI governance reviewed at contract renewal or incident trigger. No continuous visibility into third-party AI behaviour touching the firm's regulated perimeter.
AiRE outcome	Automated discovery and classification of AI assets across the estate, continuously updated. Full visibility into systems, data flows, and control status in real time.	Obligation mapping automated against the AI asset inventory. Gaps between regulatory requirements and current controls surfaced automatically, with prioritised remediation paths.	Structured audit trail generated at the point of execution for every AI-assisted decision. Explainability documentation produced automatically and stored in a retrievable, regulator-ready format.	Continuous monitoring of third-party AI activity against defined governance parameters. Anomalies and policy breaches flagged in real time, with escalation paths pre-configured.

Critically, each of these capabilities is delivered with change management and adoption built in, because the most technically sound governance architecture fails if the people responsible for operating it do not understand or trust it.

QuantSpark embeds a working prototype inside your existing infrastructure in two weeks. Not a proof of concept. Not a pilot in a sandbox. A functioning governance capability running against real workflows, with real data, inside your own environment and generating the evidence trail your regulators will eventually ask to see.

Section VII

A Framework for Action

PHASE 01

Map the Estate Before You Govern It

You cannot govern AI you cannot see. The first step is a complete inventory: every system in use, every data flow it touches, every third party with access, and every use case, whether formally approved or not. Most organisations discover significantly more AI in active use than their governance frameworks account for. That discovery is uncomfortable. It is also the only honest starting point.

PHASE 02

Redesign Governance Into the Architecture, Not Onto It

The most expensive governance failure is late-stage governance: controls added after systems have been built, pipelines have been committed, and commercial obligations have been entered into. Governance designed into the architecture from the start is both more effective and substantially cheaper. The question to ask at every design review is not “does this comply?” but “how does compliance get enforced at runtime?”

PHASE 03

Build the Orchestration Layer

Orchestration is the operational translation of governance policy into system behaviour. It connects model access, data permissions, approval workflows, override logging, and audit trails into one managed environment. Deployed inside the organisation’s own infrastructure, not as an external SaaS layer, orchestration gives compliance teams continuous visibility and control without creating new data sovereignty risks.

PHASE 04

Extend Governance to the Ecosystem

Internal AI is the part you control. Third-party AI is the part that is most likely to produce a regulatory event you did not anticipate. Phase four extends the orchestration layer to the supplier and partner ecosystem by defining what third-party AI can and cannot do with data that touches the firm’s regulated perimeter, and building the monitoring infrastructure to enforce it.

PHASE 05

Train for Governance Fluency, Not Just Tool Adoption

Governance architecture is only as effective as the people operating it. Legal, risk, compliance, technology, and business teams need shared fluency in what the controls do, how to use them, and what to escalate. The difference between governance on paper and governance in practice is almost always a training gap, not a technology gap.

PHASE 06

Monitor Continuously and Improve Systematically

Governance is not a project with a completion date. Models drift, integrations change, regulatory obligations evolve, and new use cases are added continuously. Phase six establishes the feedback loop: regular monitoring against defined governance metrics, systematic identification of drift and exceptions, and a structured process for updating controls as the AI estate changes.

Conclusion

The organisations that govern AI well do not treat governance as a capability.

That reframe matters because it changes what gets built and when. Compliance as a checklist produces controls that sit outside the system. Orchestration as an operating model produces controls that are the system, embedded in the architecture, running continuously, and generating the evidence trail that allows the organisation to move faster, not slower, because its AI is demonstrably trustworthy.

The regulatory environment will continue to harden. The EU AI Act's obligations are expanding, UK regulators are sharpening their supervisory expectations, and the organisations already operating with orchestration-led governance are pulling further ahead. Not just in their ability to demonstrate compliance, but in their ability to deploy AI at scale in the first place.

The question for every compliance, risk, and technology leader reading this is whether to act now, while the control architecture can be built on your terms, or later, when it must be built under scrutiny. Not acting is out of the question.

In a regulated world, the future belongs to businesses that can prove not only that their AI is powerful, but that it is governable.

AiRE exists for exactly this moment.

To book a free discovery session or discuss AiRE deployment in your organisation:

For more information

Visit www.launchaire.com

Email contact@quantspark.com

Selected sources

[...]

QuantSpark

Authors

Adam Hadley CBE

CEO & Founder
adam.hadley@quantspark.com

Wanda Haddock

Head of AI Transformation
wanda.haddock@quantspark.com

Contact

www.quantspark.ai

contact@quantspark.com

230 Blackfriars Road
London
SE1 8NW

MK417