Ø Occamly

WHITEPAPER · 2026

The Pragmatic AI Governance Operating Model for UK Regulated Industries

Occamly · Jake Stennett — independent advisory in AI, enterprise and solution architecture.

occamly.com · "The simplest architecture that works."

Executive summary

Most AI governance in regulated organisations fails the same way: it produces documents, not control. The fix is not more policy. It is a lean operating model that sits inside delivery.

- **Governance only works when it ships with the work.** If a review gate sits alongside delivery as a separate committee, it gets routed around. Put the gate inside the delivery path and it cannot be skipped.
- **Tier the risk, then tier the controls.** Not every AI use case needs the same scrutiny. A reporting summariser and a customer-facing eligibility model are not the same risk. Treat them the same and you either over-control the trivial or under-control the dangerous.
- **An AI register is the spine, not the paperwork.** One list of every AI use case, its tier, its owner, and its status. Without it you cannot answer the board's only real question: what AI are we running, and who is accountable for it.
- **Stand it up lean in 90 days, not as a 12-month programme.** Intake, classification, a register, and one working gate first. The rest can wait. A perfect framework that lands in month 9 governs nothing in months 1 to 8.
- **Post-Omnibus, the EU AI Act high-risk obligations are deferred, not cancelled.** This is a window to build the capability calmly, not a reason to stop. Prioritise on real risk and customer harm, not on a deadline that has moved.

The problem: why AI governance becomes shelfware

The pattern is consistent across regulated organisations. A governance framework gets commissioned. It arrives as a long document. It is thorough, defensible, and ignored.

3 things cause this.

It is built as a document, not a process. A 90-page framework describes controls. It does not enforce them. Delivery teams read it once, file it, and carry on. The framework has no hook into how work actually moves from idea to production.

It lives in the wrong place. Governance committees that meet monthly, separate from delivery, become a queue people learn to bypass. If shipping is faster without the gate, the gate loses. In regulated sectors the people building AI are usually under delivery pressure. Governance that adds friction without sitting in the critical path gets routed around.

Nobody owns the controls. A control with no named owner is a sentence in a document. When the auditor asks who signs off model risk, the answer is a shrug and a re-org. Ownership is the difference between a control and a wish.

The result is governance theatre: the appearance of control, with none of the substance. It survives audit on paper and fails the first time a model does something the organisation cannot explain.

The operating model

The model has 7 components. Each is deliberately minimal. The discipline is keeping them small.

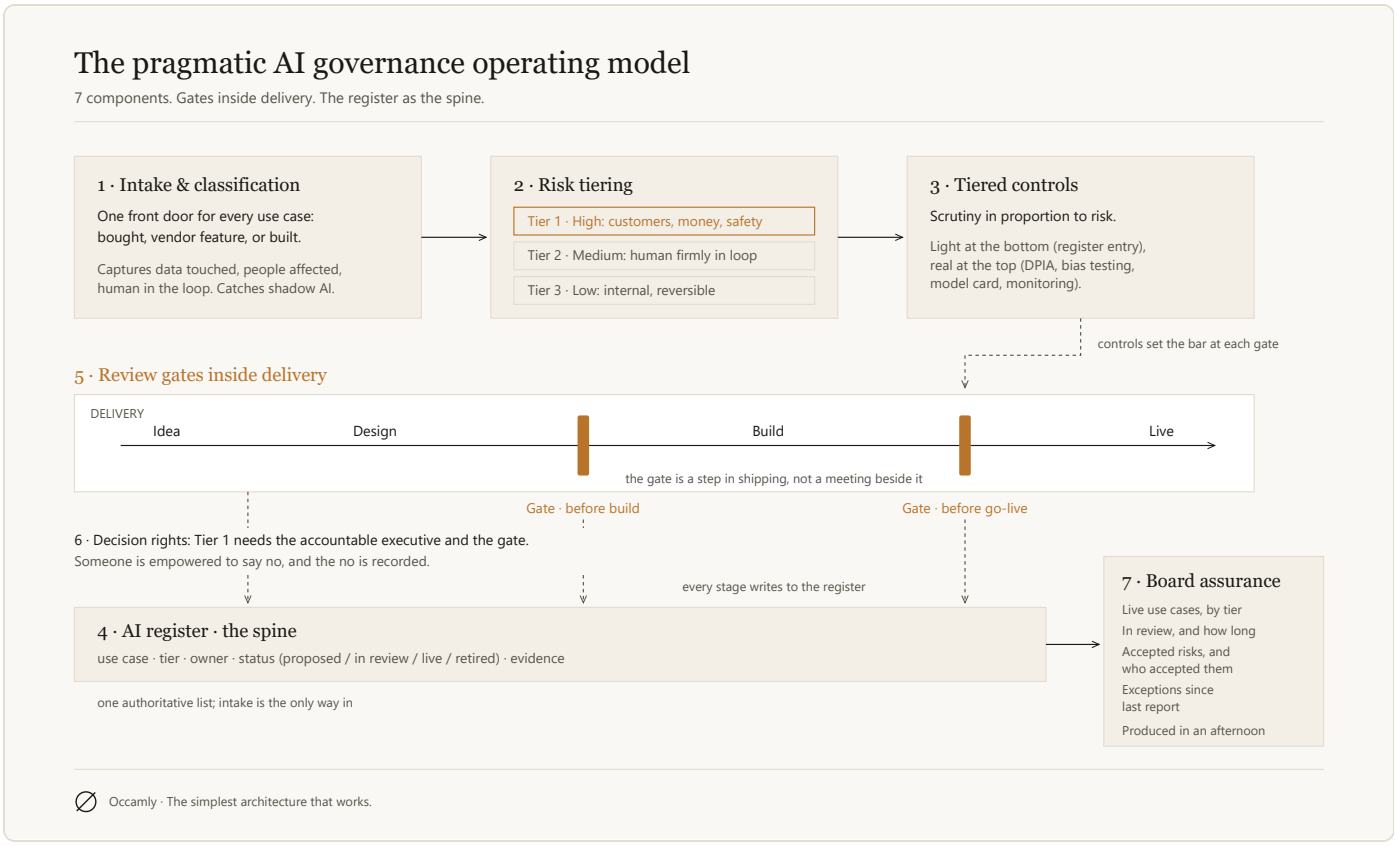

The operating model at a glance: one intake, risk-tiered, controls by tier, review gates inside delivery, a register as the spine, and board assurance drawn from it.

1. Use-case intake and classification

One front door. Every AI use case, whether a bought tool, a vendor feature, or something built in-house, enters through the same intake. The intake captures the basics: what it does, what data it touches, who it affects, and whether a human is in the loop on the decision.

This is also where you catch shadow AI: tools adopted inside a team without anyone classifying the risk. If there is no front door, you do not have a register. You have a guess.

2. Risk tiering

Classify each use case into a small number of tiers; 3 is usually enough.

Tier	Description	Examples (generic)
Tier 1: High	Affects customers, safety, eligibility, money, or legal rights. Limited or no human override.	Automated eligibility decisions, customer-facing decisions with consequence, anything safety-adjacent.
Tier 2: Medium	Material internal impact, or customer-facing with a human firmly in the loop.	Drafting assistance reviewed before it goes out, internal prioritisation tools, decision support.
Tier 3: Low	Internal productivity, low consequence, easily reversible.	Meeting summaries, internal search, code assistance with review.

The tier is a decision, not a calculation. Keep the criteria written down and keep them blunt. If tiering takes a workshop, it is too complex to use.

3. Tiered controls

Match the control to the tier. This is the core of pragmatic governance: scrutiny in proportion to risk.

Control area	Tier 3 (Low)	Tier 2 (Medium)	Tier 1 (High)
Approval	Self-service / register entry	Named owner sign-off	Formal gate review + accountable exec
Data check	Light DPIA screen	DPIA	DPIA + data lineage evidence
Human oversight	Optional	Human-in-the-loop required	Human-in-the-loop + documented override path
Testing/eval	Basic	Pre-deployment evaluation	Evaluation + bias/robustness testing + monitoring
Documentation	Register entry	Use-case record	Full model card + audit trail
Review cadence	Annual	Periodic	Continuous monitoring + scheduled review

The point is what Tier 3 does *not* require. If a meeting summariser needs the same paperwork as a customer eligibility model, people will avoid the register entirely, and you lose visibility of everything. Lean controls at the bottom buy you compliance at the top.

4. The AI register

One authoritative list. Every use case, its tier, its owner, its status (proposed / in review / live / retired), and a link to its evidence. Nothing more elaborate is needed to start.

The register is the single artefact that answers the board and the regulator. It is also the thing that decays fastest if intake is not the only way in. Protect the front door and the register stays true.

5. Review gates inside delivery

This is the component that separates governance that works from governance that does not.

Do not build a parallel governance committee that delivery has to visit. Put the gate inside the delivery process itself, at the points where work already pauses: before build, before go-live. The gate is a defined step in the path to production, not a separate meeting that competes with it.

Concretely: the use case cannot move to its next delivery stage until the tier-appropriate control is satisfied. For Tier 3 that may be an automatic check. For Tier 1 it is a real review with a real decision. The gate is mandatory because it is structurally part of shipping, not because a policy says it should be.

6. Decision rights

Write down who decides what. Ambiguity here is where governance quietly dies.

Decision	Who decides
What tier a use case is	Governance function (with appeal route)
Approve a Tier 3 use case	Use-case owner
Approve a Tier 2 use case	Named accountable owner
Approve a Tier 1 use case	Accountable executive + governance gate
Override / accept a risk	Named risk owner, recorded
Retire or pause a live use case	Governance function + owner

The most useful sentence in any AI governance model is "sometimes the answer is no." Make sure someone is actually empowered to say it, and that the no is recorded.

7. Assurance and evidence for boards

Boards do not want the framework. They want assurance. Give them a short, regular view drawn straight from the register:

- How many AI use cases are live, by tier.
- How many are in review, and how long they have been there.
- Any Tier 1 use case operating with an accepted risk, and who accepted it.
- Exceptions and overrides since the last report.

If you cannot produce that from your governance setup in an afternoon, the setup is not generating evidence. It is generating documents. Those are not the same thing.

Worked example: one use case through the model

A fictional case, end to end. Northgate Utilities, a mid-size UK water company, wants a model that flags customers for its priority services register, the list of households needing extra support during supply interruptions.

Intake. The team submits through the front door. The intake records what the model does (recommends customers for priority support), the data it touches (billing, contact history, third-party vulnerability data), who it affects (customers), and the human in the loop (a case officer confirms every flag).

Tier. Tier 1. It affects eligibility for support, it touches vulnerable customers, and a wrong "no" causes real harm. The case officer lowers the risk; they do not lower the tier.

Controls. The Tier 1 set applies: DPIA with data lineage evidence, a documented override path for case officers, bias testing across customer segments, a model card, and live monitoring.

Register. An entry is created at intake: use case, Tier 1, named owner (Head of Customer Operations), status "in review", links to the evidence.

Gate. At the before-go-live gate, the review finds the override path is not documented: case officers can overturn a flag, but nothing records why. The gate returns it. Two weeks later the override log exists, the accountable executive signs off, and the status moves to "live".

Board view. One line in the next assurance report: "Priority services flagging. Tier 1. Live. Owner: Head of Customer Operations. No accepted risks. 1 gate return, resolved."

Nothing exotic happened. The use case shipped 2 weeks later than planned, with a gap fixed that would otherwise have surfaced in front of a regulator. The gate found it, the register recorded it, and the board can see both. That is the whole model working.

Standing it up in 90 days

Lean-first. Build the smallest thing that controls real risk, then add. The sequence below front-loads the components that give you visibility and stops the things that produce paperwork without control.

Days 0–30: see what you have

- Define the 3 tiers and the tiering criteria. One page.
- Stand up the AI register. A spreadsheet is acceptable for week one. Tooling can come later.
- Run a discovery sweep: find the AI already in use, including vendor features and shadow tools. Populate the register with what exists.
- Name an accountable owner for each existing use case.

Output: you now know what AI you are running and who owns it. Most organisations cannot answer this. That alone is worth the month.

Days 30–60: control the front door

- Build the intake. One front door for every new use case.
- Wire the first review gate into delivery for Tier 1 use cases only. Start where the risk is, not everywhere at once.
- Write the decision-rights table and get it signed off.
- Draft the tier-appropriate control set, but only fully specify Tier 1. Tier 2 and 3 can be lightweight for now.

Output: no new high-risk AI reaches production without a real review, and everything new is logged.

Days 60–90: prove it and report

- Produce the first board assurance view from the register.
- Tighten Tier 2 controls based on what you learned running Tier 1.
- Add monitoring for live Tier 1 use cases.
- Document the model itself in a few pages, not 90.

Output: a working, evidenced governance capability the board can see, in a quarter.

What waits. Full tooling, exhaustive Tier 3 controls, a polished policy library, an org-wide training programme. All useful. None of it controls risk in week one. Build them once the spine is load-bearing.

Anti-patterns

Recognisable failure modes. If you see these, the model has drifted.

The 90-page report that ends in "it depends." Thoroughness as a substitute for a decision. A governance artefact that does not tell someone what to do on Monday is not governance. It is research.

Governance theatre. Committees, RACIs, and policy documents that produce the look of control and none of the function. The tell: lots of meetings, no record of anything ever being stopped or changed.

Controls nobody owns. Every control needs a named human. "The team" is not an owner. If you cannot name who signs off model risk, you do not have that control.

Tiering everything as high. Caution that collapses into paralysis. If everything is Tier 1, nothing is, and delivery learns to avoid the register entirely.

The parallel committee. Governance that runs alongside delivery instead of inside it. It becomes a queue, then a bottleneck, then a thing people route around. Put the gate in the delivery path or accept that it will be skipped.

Register drift. The register exists but is not the only way in, so it slowly stops reflecting reality. A register that is 70% accurate is more dangerous than no register, because people trust it.

The post-Omnibus reality

The regulatory picture changed, and the change rewards calm over panic.

The EU AI Act came into force on 1 August 2024. Transparency duties apply from 2 August 2026. The high-risk obligations, the demanding part for most regulated organisations, were deferred by the Digital Omnibus: 2 December 2027 for Annex III high-risk systems and 2 August 2028 for Annex I. The Omnibus was provisionally agreed in May 2026; formal adoption is expected before 2 August 2026, so the dates may still shift.

Whether it binds a UK organisation depends on EU exposure. The Act reaches beyond the EU. A UK organisation is in scope where it places an AI system on the EU market, or where the output of its AI is used in the EU, for example EU customers, EU operations, or EU applicants screened by a model. Purely UK operations, with no EU market and no EU use of outputs, sit outside it. The UK itself remains sector-led: there is no equivalent cross-sector AI statute in force or imminent, and the headline 2026 measure, the Regulating for Growth Bill, is about regulatory sandboxes, not an EU-style rulebook. The practical position for most UK regulated organisations is two-track: meet the EU Act where you touch the EU, and meet existing UK sector obligations everywhere else.

Do not read the deferral as a reprieve to do nothing. Read it as a window.

What it means for prioritisation:

- **Build on real risk, not the calendar.** The high-risk deadline has moved, but customer harm has not. Prioritise the use cases that can actually hurt someone or the organisation, not the ones a deadline once forced to the top.
- **2 August 2026 is not a high-risk hard deadline.** Transparency duties apply from then. The heavy high-risk conformity obligations do not. Anyone selling you a 2026 high-risk panic is selling urgency, not accuracy.
- **Use the runway to build capability, not paperwork.** A lean operating model running well by the time obligations bite is worth far more than a documented framework written to a deadline and never operationalised.

The honest framing is "navigating the confusion." The dates have moved and may move again. The right response to a moving deadline is a capability that is genuinely embedded, because that holds whatever the date turns out to be.

One-page checklist and maturity ladder

Checklist: do you actually have control?

- [ ] One front door: every AI use case enters through a single intake.
- [ ] A live AI register that reflects reality, including vendor and shadow AI.
- [ ] Every use case has a named, accountable owner.
- [ ] 3 risk tiers with blunt, written criteria.
- [ ] Controls scaled to tier (light at the bottom, real at the top).
- [ ] At least one review gate sitting inside delivery, not alongside it.
- [ ] A written decision-rights table, including who can say no.
- [ ] A board assurance view you can produce in an afternoon from the register.
- [ ] Evidence of at least one use case being changed, paused, or stopped by the model.

If you cannot tick the last box, you have process, not governance.

Maturity ladder

Level	State	What it looks like
0: Blind	No register, no intake.	AI in use across the org, no one can list it. Shadow AI unmanaged.
1: Visible	Register exists; use cases owned.	You can answer "what AI are we running and who owns it."
2: Tiered	Risk tiers and proportionate controls applied.	Scrutiny matches risk. The trivial is light; the dangerous is reviewed.
3: Embedded	Gates sit inside delivery; decisions recorded.	Nothing high-risk ships without review. The model cannot be routed around.
4: Assured	Board-level evidence produced routinely; model self-corrects.	Governance generates assurance, monitors live use cases, and stops things when needed.

Most regulated organisations are at Level 0 or 1 and believe they are at Level 3 because they have a framework document. The ladder is climbed by operating, not by writing.

Working with Occamly

Occamly is independent advisory in AI, enterprise and solution architecture. The brand line is the method: the simplest architecture that works.

If your AI governance is a document nobody uses, or you need to stand up a capability that survives delivery rather than sitting beside it, that is the work. The approach is conclusion-first, evidence-led, and willing to tell you when the answer is "don't."

To discuss an engagement, enquire by email via **occamly.com**. No bench, no associates, no hard sell.
