The AI Control Manifesto

Untracked models, rogue agents, $20M fines. AI use explodes while CROs are kept in the dark.
governr is your AI Control Room.

Thushan Kumaraswamy
Co-Founder governr
March 2026

What We Believe

1. We believe AI proliferation across financial services is inevitable and accelerating.

AI proliferation is inevitable. The $97B market by 2027 shows 88% adoption in trading and fraud detection. Models, agents, and APIs multiply across functions, demanding robust automated governance at enterprise scale.

2. We believe AI agents are digital workers requiring the same controls as human employees.

AI agents are digital workers requiring controls like human employees. Agents need access rights, reporting structures, and performance monitoring. Without proper onboarding and revocation protocols, liability compounds across the 95% of interactions agents now handle.

3. We believe manual governance collapses beyond hundreds of assets.

Manual governance collapses beyond hundreds of assets. Spreadsheets fail at scale; quarterly audits miss drift. With 71% of firms below 30% governance coverage, automated oversight is the only viable path to regulatory compliance.

4. We believe senior executives bear personal accountability to demonstrate enterprise-wide AI control.

CROs face $20M fines and disqualification. EU AI Act, FINRA, and Bank of England demand live traceability. Quantified exposure reports must replace vague assurances to satisfy emerging regulations.

5. We believe effective control requires integrated, real-time automation.

Static processes and periodic reviews cannot contain AI's exponential velocity. Real-time oversight across every model, agent, and API is not optional; it's the essential discipline required to harness AI's power while preserving operational integrity and regulatory trust.

The Explosion Accelerating Across Financial Services

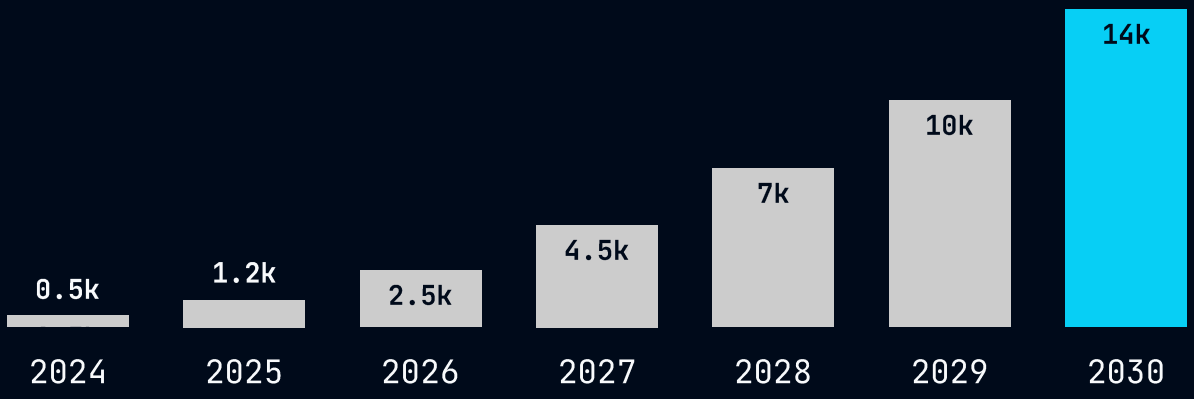

FIG 1: Projection of the average number of AI assets deployed per financial services firm, based on current estimates from JP Morgan and the Bank of England

Financial services will commit $97 billion to AI by 2027, with 88% adoption across fraud detection, trading signals, and client servicing where chatbots handle 95% of interactions. Every department deploys models, agents, and APIs weekly, while generative AI scales toward $17.9 billion by 2035. Front office algorithms, middle office reconciliation, and back office compliance all run on AI today.

Unfortunately, no firm governs more than 20% of these assets effectively. Shadow AI spreads through ChatGPT plugins, internal agents, and third-party APIs. Recent analysis identifies 1.5 million enterprise agents carrying rogue risk, including hallucinations recommending toxic investments and drift generating million-dollar losses.

For example, Air Canada's chatbot approved $2,100 refunds it couldn't deliver, creating legal liability in hours.

These digital workers require employee-grade controls: defined access rights, reporting lines, daily KPIs, and termination protocols. AI agents currently operate without onboarding registries or performance monitoring, creating PII exposure and SEC violations that compound hourly.

When the board demands answers—"How many agents handle client money? What risks materialised last week?"—model risk teams cannot respond. This transparency gap has created a crisis. The $378M AI governance market grows at 38% CAGR because proliferation has outpaced every manual process.

Why Model Risk Teams Cannot Govern AI Velocity

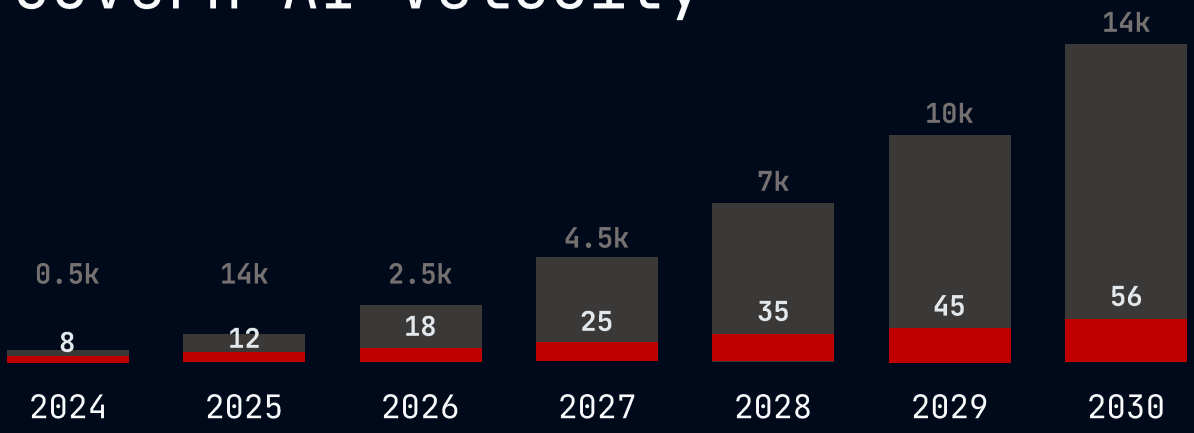

FIG 2: Projection of the average number of model risk staff in financial services firms, based on 15-20% rise annually (PWC)

Model risk functions attempt to govern dozens of models through established SR 11-7 quarterly validation cycles, static model cards, and Excel inventories designed for 2015 machine learning pilots. These tools were never built for 2026 agent swarms. While headcount grows annually, training lags six months behind, and spreadsheets reach capacity at a few hundred assets. Manual governance fails velocity completely.

AI velocity fundamentally breaks this traditional model. Agents update hourly through vendor APIs, drift evolves daily from market regime shifts, and new shadow agents emerge weekly via developer tools like GitHub Copilot. Seventy-one percent of deployments carry less than 30% governance coverage, as manual processes scale linearly while AI complexity grows exponentially.

Regulators have anticipated this gap. The EU AI Act mandates continuous traceability for high-risk systems rather than quarterly PDF reports. FINRA's suptech (supervisory technology) framework demands runtime evidence of dynamic controls. Fines of $20 million and director disqualification now target CROs unable to prove live oversight, while the Bank of England flags systemic risk from ungoverned velocity.

Quarterly audits miss 90% of live drift events, and shadow AI evades review entirely. Without automation, the Knight Capital repeat becomes routine across thousands of unchecked instances. Headcount cannot close the gap, and consultants charge $1 million to map just 5% of assets. Scale demands automation, as the $378M market grows 38% annually because manual processes cannot survive this environment.

The Five Layers of the AI Control Room

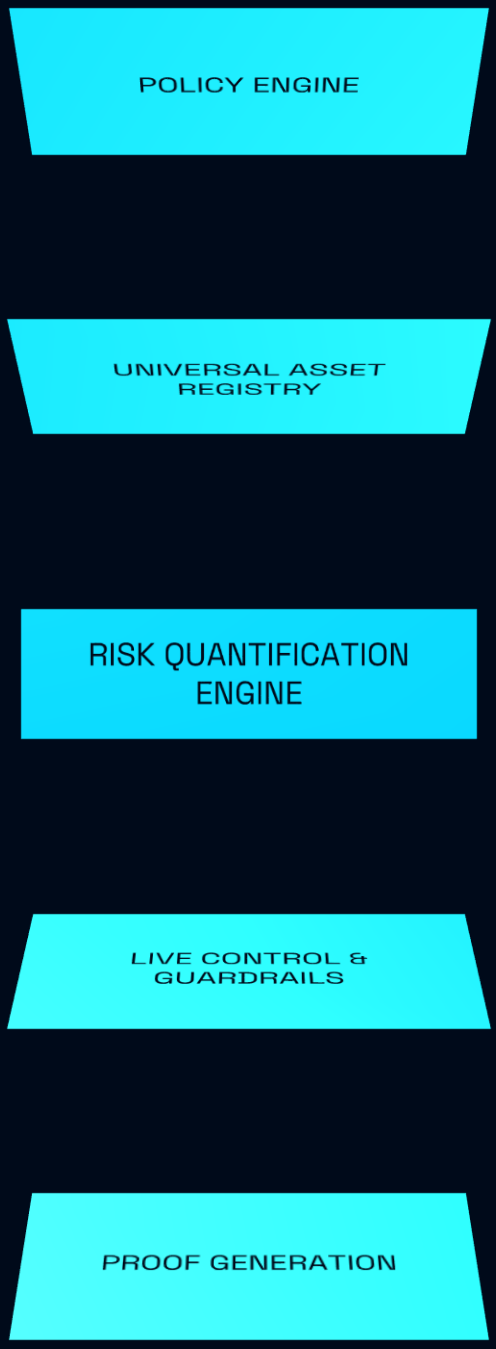

1. POLICY ENGINE

Regulatory compliance must become executable code. EU AI Act, FINRA suptech, SR 11-7, and ISO 27001 requirements exist as JSON parameters. One change cascades across thousands of assets. Business leaders set risk appetite directly with no developers required.

2. UNIVERSAL ASSET REGISTRY

AI assets live across clouds, GitHub agents, OpenAI APIs, and LangChain. Dependencies flow agent → model → dataset → vendor. Risk tags apply automatically. No asset escapes the registry's immutable truth across 10,000+ instances.

3. RISK QUANTIFICATION ENGINE

AI risk demands dollar quantification. Simulations and hourly forecasts reveal "doubling drift triples $2.3M exposure." CROs receive portfolio-level materiality, not vague probabilities.

4. LIVE CONTROLS & GUARDRAILS

Runtime violations demand instant response. AI telemetry matches behaviour against policy tags, creating tickets and CI/CD gates. Digital workers receive KPIs like human employees. Enterprise kill switches protect against rogue deployments.

5. PROOF GENERATION

Regulators demand instant proof. FINRA packs and EU AI Act documentation hyperlink directly to telemetry states and policy parameters. 60-minute SLAs deliver PDF/API exports. Every claim carries its own immutable audit trail; perpetual audit-readiness is non-negotiable.

What Does Good Look Like?

Chief Risk Officer Radar

Every AI asset appears categorised by risk tier, with violations mapped to specific business lines and compliance velocity tracked in real time. Automated remediation runs to defined service levels, so issues resolve proactively rather than crisis management.

Board Reporting

AI concentration risks become visible across trading, client interaction, and fraud systems, with vendor dependencies and worst-case scenario forecasts available. Board packs compile with one action, giving leadership complete clarity without manual assembly.

Regulatory Examinations

Suptech submissions generate automatically, pulling from asset inventories, policy mappings, and remediation evidence, all hyperlinked to source data. Model risk teams shift from documentation drudgery to strategic oversight.

Incident Management

Digital worker failures trigger containment workflows immediately, throttling, human approval gates, and developer ticketing, with before-and-after risk scoring preserved for audit. No repeat of past multimillion-dollar incidents.

Shadow AI Mitigation

Untracked agents surface through regular scans, receive automated policy tagging and ownership assignment, and convert to governed assets within days rather than months.

Achieve this standard

governr is the AI Control Room that delivers the maturity that regulators and Boards are demanding.

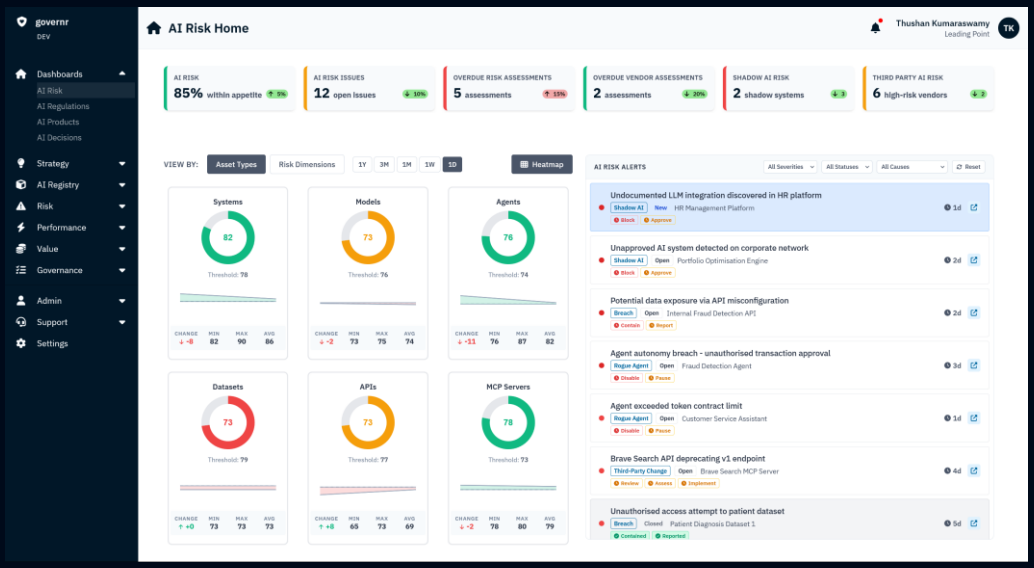

Thushan Kumaraswamy
Co-Founder & CPO
thush@governr.ai

Book your tailored demo today