governr

One Controller for All Enterprise AI

Inside the Machine: How Regulated Sector Enterprises Are Taking Control of AI

Based on 25+ interviews across regulated firms in the US and UK (FS, healthcare, builders).

Know what AI is running, what it touches, and who owns it, with the AI Control Plane standard and a 90-day KPI plan.

Executive Summary

Regulated organisations are moving from AI experimentation to AI at scale. The constraint is no longer interest or budget. It is control.

This briefing is grounded in fieldwork: 25+ interviews across regulated enterprises in the US and UK, spanning financial services, healthcare, and AI builders. The pattern was consistent: leaders want the upside of AI, but lack an operating capability that can keep pace with constant change.

> **In 60 seconds: what leaders should take away**
>
> - AI sprawl comes from employees, builders, and vendors. Cover one and you miss the risk.
> - 2026 regulatory obligations converge on provable control, not policy decks.
> - The AI Control Plane is the operating layer for daily monitoring and audit-ready evidence.
> - If you cannot detect material AI change within one business day, you do not have control.

Adoption is already widespread. A 2024 WalkMe survey found 78% of employees admit using unapproved genAI at work.[1] Gartner forecasts that by 2026, more than 80% of software vendors will embed genAI into enterprise applications.[2]

> **Field signal:**
> In a board meeting at a large US bank, a director asked: “Do we know what AI is running?” The room went quiet. That silence is the gap this report addresses.

This report defines the AI Control Plane operating standard: continuous oversight that shows what AI is running, who owns it, what it can access, what changed, and what controls apply across employees, builders, and vendors.

What follows: why traditional governance breaks at scale, what signals leaders need to regain visibility, and a Now/Next/Later plan with executive KPIs to prove whether the organisation is steering AI, or discovering it under pressure.

The Executive Decision: Treat AI Oversight as Infrastructure, Not Policy

Most regulated firms are not failing on intent. They are failing on operating capability. Policies and periodic reviews were built for systems that change slowly. AI does not. It expands through everyday behaviour, developer workflows, and vendor upgrades.

The right question is not "do we have governance?" It is operational:

- **Coverage**: Do we know where AI is used in our regulated processes?
- **Change**: Can we detect material changes inside one business day?
- **Accountability**: When something changes, can we route it to an owner immediately?
- **Evidence**: Can we produce an audit trail without a multi-week scramble?
- **Regulatory readiness:** Can we produce Tier 1 evidence within 72 hours when asked?

If you can't answer these reliably, you will discover AI under pressure: in an exam, incident, or board escalation.

What we heard:

- A risk leader at a major regulated institution told us they are trying to enable the business responsibly, not be a blocker, but their oversight cadence and tooling cannot keep pace with AI change.

- **A senior leader at a large cloud platform provider told us that when controls are clear and friction is removed, adoption accelerates and commercial impact shows up quickly.**

- A former head of AI risk at a regulated financial services firm told us that building the inventory and control framework once allows organisations to meet multiple regulatory expectations at once, instead of recreating a new compliance project for each new requirement.

The Decision: Which Sprawl Problem Are You Actually Governing?

AI governance breaks at scale for the same reason operational resilience breaks: change outpaces controls built on periodic reviews and manual updates.

AI enters through three feeders. Each requires different signals. The failure mode is building oversight for one and assuming it covers the rest.

1. Vendor AI Expands Through Upgrades

AI increasingly enters as routine product updates, not new purchases. Annual reviews capture intent, but miss the operating surface once features ship and configurations expand. **Gartner forecasts 80%+ of ISVs will embed genAI by 2026.**[3]

2. Shadow AI Scales Faster Than Policy

Unmanaged accounts and unsanctioned tools are now a default way employees draft, search, analyse, and code. At enterprise scale, policy-only control fails because usage spreads faster than review processes can detect. **WalkMe found 78% of employees use unapproved genAI.**[4]

3. Development Shifts Risk Upstream

The highest-impact choices are made during experimentation, when architectures become dependencies and toolchains harden. That is also where ownership, documentation, and monitoring are weakest. **McKinsey reports 88% using AI in at least one function, and 62% experimenting with AI agents.**[5]

What we heard:

- A global buy-side firm told us they encouraged AI for productivity, then hit the predictable wall: visibility collapsed. Leadership couldn't reliably track which tools were used or where data was flowing.
- An analytics firm described a recent engagement where roughly 300 citizen developers were building AI-powered utilities outside formal governance, creating a parallel AI estate that risk teams could not see.

Do You Govern Models, or Do You Govern Workflows?

These are the three reasons the control cycle breaks once sprawl exists

As AI becomes agentic, the unit of risk shifts. It is no longer just model quality. It is workflows, permissions, connectors, and change velocity. If oversight runs quarterly and AI changes daily, control is assumed, not proven.

4. Agents Shift Oversight to Workflows

Oversight cannot focus only on individual models. The risk increasingly sits in agentic workflows that take actions, chain tools, and expand data access through integrations and permissions. This shift is already underway in production.

5. Controls Decay Between Review Cycles

When systems change daily, assurance decays inside the review cycle. What was approved last quarter is often not what is running this week due to updates, prompt changes, workflow edits, and permission creep. Continuous, daily oversight becomes the only viable operating standard.

6. Inventory Becomes a Hard Requirement

Inventory is shifting from optional to required. Regulations and standards increasingly expect evidence of what AI is in use, where it runs, who owns it, and what data it touches. If you cannot maintain a living inventory of what is in production, you cannot credibly govern it.

The Category Shift: From Periodic Review to Operational Control

Committees and policies are necessary. They are not control. When AI enters via vendor switches, citizen builds, and personal accounts, oversight must start from telemetry and change detection, not static lists.

The AI Control Plane is the missing layer between strategy and defensible control. It is not a committee. It is an operating capability that runs every day: it discovers what is in use, routes material change to accountable owners, and keeps controls and evidence current.

The decision is straightforward: treat AI oversight as infrastructure (always-on), or accept periodic oversight and discover change under pressure. The firms moving fastest are not writing more policy. They are building operating control.

What we heard:

- A large US bank planning aggressive AI deployment in 2026 told us the real problem is not intent, it is churn: “Every vendor is going to ship another AI switch.” Their current risk intake was described as “basically an Excel form” - and they admitted it cannot keep up.
- A global buy-side firm told us they were tracking more than ten AI-specific regulatory obligations landing in January 2026 alone, spanning EU requirements and multiple UK and US regulators. The compliance load is compounding, while the underlying inventory is still incomplete.
- A former COO from a hedge fund put it in historical terms: “This looks like the early days of model validation, except the gap is wider, the change rate is higher, and the consequences compound faster. Same gap - now exponentially worse with AI.”

2026 Readiness: Can You Prove AI Control on Demand?

Across 2026, AI obligations move from emerging guidance to operational expectations. The EU AI Act is the clearest signal, with staged applicability and broader obligations coming into effect by August 2026. In the US, federal direction has shifted through executive action and agency policy, while OMB M-24-10 formalised requirements for federal agencies' use and oversight of AI.

The point is not that every regime is identical. It's that expectations converge on the same operating proof: a current inventory of AI in use (including third-party and embedded), accountable ownership, evidence that stays current as systems change, and demonstrable control over permissions, data access, and vendor changes.

This is where periodic review fails. If you can't produce current evidence quickly, you don't have a compliance gap, you have a control gap.

An AI Control Plane is the advantage: daily oversight that keeps inventory, accountability, and evidence current, so you can respond to new obligations without rebuilding a new compliance project each cycle.

What we heard:

- A CTO at a regulated insurer described the "multi-vendor agent boundary" problem: agents from multiple vendors now operate in the same environment, each with its own controls, permissions, and defaults, but no vendor can govern the full workflow end to end.

The next section shows where inventories break first, and why visibility collapses before governance catches up.

The Proof Gap: Adoption Is Outpacing Evidence

Even the most mature sectors have model inventories, not AI inventories. Adoption is outpacing proof. Financial services leads on model governance, yet the evidence shows a widening gap between what firms run and what they can evidence.

If you can't produce current AI evidence quickly, you don't have a reporting gap. You have a control gap.

Benchmark	Finding	Implication
KPMG	83% of banks have a formal model inventory, but only 17% include AI/ML/AML/fraud/ESG models.[6]	Model inventories are not AI inventories. Coverage is thin where risk is now emerging.
BoE/FCA	75% of firms already use AI. 33% of AI use cases are delivered by third parties, and 46% report only partial understanding of the AI they use.[7]	Third-party delivery expands the AI surface area faster than firms can see or evidence.
Netskope	95% of FS organisations now use genAI, with an average of 10 different genAI apps. Usage includes regulated data, IP, and source code.[8]	AI usage is processing regulated data outside traditional model risk controls and inventories.
Black Book	Just 22% of surveyed hospital leaders were confident they could produce a complete AI audit trail within 30 days. Barriers: limited documentation (41%), unclear ownership (33%).[9]	Healthcare is in the same position as financial services: audit-readiness lags AI adoption.
OMB	2024 federal inventory counted 1,700+ AI use cases. After Dec 2024, systems without required safeguards must be stopped.[10]	Inventory is now treated as enforceable control. Expect this standard to migrate into regulated audits.

What Happens When Inventory Is Incomplete

The consequences of incomplete AI inventory are not abstract “ethics risks”. They are operational, regulatory, and security outcomes that leaders can expect to face:

- Delayed launches due to late discovery of ungoverned AI
- Examination friction and remediation projects when regulators ask questions you cannot answer
- Incident response failures when permissions have drifted beyond what was approved
- Vendor renegotiations and procurement delays when AI features appear mid-contract
- Insurance claim denials when governance evidence is incomplete or stale

What leaders lose when this is missing:

- **Business cost:** delayed launches, procurement renegotiations, insurance denials, remediation programs.
- **Regulatory cost:** exam friction because you can’t evidence what’s running, what changed, and who approved it.

What we heard:

- A former COO at a hedge fund reframed the leakage risk: “We built firewalls to stop threats coming in. The AI era requires firewalls to stop proprietary information leaking out.”

- A CISO at a multi-trillion-dollar financial institution: security leaders are “drowning in the velocity of AI adoption.” The goal is not to block AI. It is to “unleash it while making it safe.”

What this means for executives:

- Mature sectors still struggle with scope and speed-to-proof.
- Daily oversight is becoming the cost of scaling AI in regulated processes.
- Executive test: Can you produce evidence for Tier 1 AI assets in under 72 hours?

What You Need to Detect: Three Feeders, Different Signals

AI enters through three feeders, and each requires different telemetry. The enterprise needs one truth.

Feeder	What to Detect	Signals That Reveal It
Employees (Shadow AI)	Which AI tools are in use, by whom, how often. Whether “approved” tools are used in unapproved ways.	SSO/identity logs. CASB discovery. DLP events flagging sensitive data movement. App telemetry surfacing usage outside SSO.
Builders (Developer Layer)	What models, agents, prompts, workflows actually exist. When behaviour changes because code or prompts changed. New data access (connectors, permissions).	Code repos: agent builds, AI projects, prompt changes. Model registries: versions, owners. Prompt/config repos. Agent logs: tool and connector usage, access patterns.
Third-Party AI (Vendor-Embedded)	When a vendor expands AI functionality. When default settings enable data sharing. When new sub-processors or models are introduced. Whether vendors can produce evidence.	Contracts/TPRM: AI disclosures, audit rights, evidence SLAs. Admin settings: AI enablement, data retention. Release notes: scope expansion. Change notices enforcement.

What this means for executives:

- Oversight that covers only one feeder is not oversight.
- Each feeder needs different signals, but the enterprise needs one truth.
- If you cannot detect privilege drift in one day, you will discover it during an incident.

Operating Model: The Control Plane as a Managed System

This is the minimum operating model that lets you prove control continuously.

1. Ownership That Matches the Operating Reality

- **Board / Risk Committee:** appetite, thresholds, dashboard requirements
- **CEO / COO**: enterprise posture, regulated-process coverage
- **CTO / CIO:** telemetry standards, integration patterns, change triggers
- **CISO**: identity, DLP, audit logging, retention
- **Risk / Compliance / Regulatory:** tiering, control set definitions, assurance method
- **Procurement / TPRM:** AI disclosure clauses, change SLAs, vendor evidence requirements
- **Control Plane Office:** system of record, triage, exceptions, reporting cadence

2. Intake Becomes a Routing Mechanism, Not a Bottleneck

- Single intake channel for new AI use, vendor AI enablement, agent deployment, new connectors, and material workflow changes.
- **Tiered routing:** low-risk fast path, high-risk deeper review, explicit exception path.

3. Control Mapping Must Be Operational

For each asset or workflow: owner, purpose, data classes touched, autonomy level, dependencies, required controls, monitoring plan, evidence artifacts, exception status.

Operating Cadence: How the Control Plane Runs Day to Day

4. Monitoring Cadence: Daily by Default

- **Always-on signals**: usage patterns, data movement, agent tool access changes, new integrations, vendor AI feature enablement.
- **Daily triage (non-negotiable)**: review detected changes, classify materiality, route to owner, apply tiered controls or log exception.
- **Weekly assurance review**: confirm tier 1 and tier 2 still meet controls; close loop on exceptions. Exceptions without expiry are policy debt.
- **Quarterly executive reporting**: coverage, drift, third-party change performance, audit readiness.

Most programs fail here because they treat change as an exception. The Control Plane assumes change is constant.

5. Exception Handling and Stop-Use

- Exceptions require: owner, compensating controls, expiry date, and revalidation schedule.
- Define stop-use criteria and rollback plans for tier 1 assets. Stop-use requires pre-agreed rollback paths for Tier 1 systems.

90-Day Path to Defensible AI Control: NOW (0–30 Days)

A phased plan with executive KPIs at each stage. Each phase builds on the last.

NOW (0–30 Days)

Establish Visibility and Ownership

- Set scope and accountability across Employees, Builders, and Vendors (including embedded AI features).
  KPI: % of business units with named AI owner (target: 100%)
- Stand up the system of record with mandatory metadata (owner, purpose, data classes, integrations, model or vendor).
  KPI: % of in-scope AI assets with complete metadata (target: 80%)
- Baseline AI usage and sprawl (sanctioned vs unsanctioned, managed vs personal accounts).
  KPI: Discovery coverage - % of AI tools in use captured on record
- Lock down regulated data paths with an approved tool strategy and managed identity.
  KPI: % of regulated users on managed accounts for approved AI
- Turn on change detection triggers for vendor AI enablement, new agents, new connectors, and permission changes.
  KPI: % of critical systems with monitoring enabled (target: 100%)
- Update procurement and renewals with AI disclosure and evidence requirements (what changed, where it runs, what data it touches).
  KPI: % of renewals with AI disclosure clauses
- Run an executive drill: answer “What AI is running?” for the top 10 regulated processes.
  KPI: % of top 10 processes with a verified AI inventory and owner
- **Now outputs (what leaders should have in-hand):** System of record, baseline sprawl report, change triggers active, top 10 process drill completed.

90-Day Path to Defensible AI Control: NEXT (30–90 Days) and LATER (90–180 Days)

NEXT (30–90 Days)

Activate Daily Monitoring and Tiered Controls

- Implement daily monitoring with alerts routed to the right owner and an executive exception summary.
  KPI: Median time to detect material AI change (target: 1 day)
- Establish a standard control path for changes: approve, document, test, and evidence.
  KPI: % of material changes with recorded approval and evidence
- Harden third-party AI oversight: map critical vendors to regulated processes, require change notices, validate defaults.
  KPI: % of critical vendors with verified AI data handling settings
- Expand beyond the top 10 processes to the top risk areas by data sensitivity and customer impact.
  KPI: % of regulated processes covered by the control plane
- **Next outputs:** Live oversight internally and vendor oversight standards enforced, coverage expanding beyond the top 10.

LATER (90–180 Days)

Move from Discovery to Assurance

- Move from discovery to assurance: continuous evidence packs for audits and board reporting.
  KPI: Evidence freshness (days since verification) by critical process
- Track permission and connector drift across agents, apps, etc.
  KPI: # of high-risk drift events per month and time to remediate
- Institutionalize governance into operating rhythm (monthly exec review, quarterly control testing, continuous monitoring).
  KPI: % of control tests passed, with documented remediation
- **Later outputs:** Audit-ready evidence on demand, measurable reduction in drift, board-level reporting that reflects what's running.

Conclusion

AI is now a daily-change system. Oversight that runs quarterly cannot control systems that change daily. Shadow AI scales as a behavior pattern, developer experimentation becomes dependency, and vendor-embedded AI turns procurement lists into partial truth.

The differentiator for regulated firms will be the AI Control Plane: an always-on layer that unifies discovery across employees, builders, and vendors, assigns accountability, tiers risk, maps controls to operating reality, and keeps assurance current as a byproduct.

What you can do now: If this report resonated, here are three concrete actions - each takes less than an hour and reveals where your organisation stands:

1. Run the Executive Drill

Pick your top 10 regulated processes. For each, answer: Is AI involved? Who owns it? When was the last change? Can you produce evidence? If the answers take longer than a day to assemble, you have a visibility gap.

2. Ask Your Top 5 Vendors One Question

Ask: “Have you introduced or expanded any AI capabilities in the last 12 months?” If the answer is yes, and it was not flagged through your risk intake, your third-party AI surface area is larger than you think.

3. Book a 30-Minute Discovery Call with governr

We can show you: what automated AI discovery looks like across code repos, vendor tools, and employee behaviour; how 60+ factor risk scoring works in practice; and what your first 30 days of continuous oversight would produce. No commitment, no lengthy procurement process. Just clarity on where you stand.

You can’t control what you can’t see.

governr inventories your AI, scores risk by context, and keeps evidence current across internal and third-party tools, including agents.

Book an executive review: governr.ai

References

[1, 4] SAP (WalkMe) (2025). Shadow AI: unapproved use and training gaps.
[2] Gartner (2024). LLM tools will drive API demand growth.
[3] Gartner (2025). 40% of enterprise apps will use AI agents by 2026.
[5] McKinsey (2025). The State of AI.
[6] KPMG. Turning compliance into a competitive edge (PDF).
[7] Bank of England (2024). AI in UK financial services.
[8] Netskope. Cloud and Threat Report: Generative AI 2025.
[9] Black Book (2025). The State of AI-Enabled Care in U.S. Healthcare
[10] White House OMB (2024). M-24-10: Advancing governance for... use of AI.