fortayconnect

RISK-FIRST PLAYBOOK · 2026

Zero leakage data fencing & shadow AI elimination.

A risk-first playbook for AI pilots in regulated financial services.

Most firms do not have an AI governance problem — they have an honour system. This paper sets out a practical architecture for stopping regulated data escaping into unapproved tools, removing the incentive for shadow AI, and protecting the AI budget as consumption pricing arrives.

Financial Services · CX & AI Advisory

7 sections
Honour system to ninety-day path

Authored by
Fortay Connect · July 2026

What's inside.

OVERVIEW

Route every interaction through one governed gateway.

Most financial services firms do not have an AI governance problem. They have an honour system — staff trusted to choose the right model, paste the right data and exercise the right judgement, with a policy PDF as the only safety net.

Why the honour system fails

That approach survived the pilot phase because pilots are small. It will not survive scale, consumption-based pricing, or a regulator that has made individual accountability the centrepiece of its supervision.

Two connected disciplines

This paper sets out a practical architecture for **zero leakage data fencing**, which stops regulated data escaping into unapproved AI tools, and **shadow AI elimination**, which removes the incentive for staff to use those tools in the first place.

It also explains why the same control layer that closes the compliance gap is the one that protects the AI budget when subsidised per-seat pricing gives way to consumption-based billing.

THE ARGUMENT IN ONE LINE

Route every AI interaction through one governed gateway, and data leakage, accountability and runaway token cost become one problem with one fix.

> **Three problems usually treated separately become one problem with one fix.**
>
> — FORTAY CONNECT

SECTION 01

The honour system is not a control framework.

Ask a compliance lead how their firm governs AI usage and the honest answer, more often than not, is encouragement.

Encouragement is not a control

Staff are encouraged to use the approved tool. Encouraged to select the cheaper model. Encouraged not to paste client data into a public chatbot. Encouragement is culture, and culture matters — but it is not a control. No firm would run payments, trading access or data loss prevention on encouragement, yet that is precisely how most AI estates operate in 2026.

The pilot succeeded, the governance never started

The pattern is consistent. An AI tool is rolled out, usually to engineering first, and adoption spreads sideways. Within months the firm has enthusiastic usage, genuine productivity gains and almost no visibility of what is being asked, what data is leaving, which models are being used or what any of it costs.

> **If your model selection policy is a request that users pick the cheaper option, you do not have a policy. You have a suggestion.**
>
> — FORTAY CONNECT

SECTION 02

Shadow AI: the risk you already have.

Shadow AI is any AI usage outside sanctioned tools and monitored channels — the analyst summarising a portfolio in a free chatbot, the developer pasting proprietary code into a browser extension, the team lead using an unapproved transcription service because the sanctioned route was slower.

01

The data gravity is severe

Client identities, positions, transaction histories and complaint records are exactly the material staff want summarised, drafted around and analysed — and exactly the material that must never leave a governed boundary.

02

The blast radius is regulatory, not just reputational

A leak into a consumer AI tool is a data protection event, a potential conduct event and a record-keeping failure at the same time.

03

Prohibition does not work

Firms that ban AI outright do not eliminate usage; they push it onto personal devices and accounts, where visibility drops to zero. Shadow AI is not a discipline problem — it is a product problem: staff route around sanctioned tools when sanctioned tools are worse.

THE UNCOMFORTABLE CONCLUSION

Shadow AI is evidence of demand. The fix is not a sterner policy PDF — it is a sanctioned route genuinely better than the unsanctioned one, wrapped in controls the user never has to think about.

SECTION 03

No new rules is not no rules.

The FCA has been explicit that it will not write a bespoke AI rulebook. That has been widely misread as breathing room. It is the opposite — the regulator is applying frameworks already in force: the Consumer Duty, SM&CR, and its existing expectations on governance, systems and controls.

01

Individual accountability is the enforcement mechanism

Under SM&CR, responsibility for AI systems attaches to named senior managers. When an AI-driven process produces a poor outcome, the question will not be which system failed but which senior manager owned it and what evidence of oversight exists.

02

The black box defence is dead

A firm cannot discharge its obligations by pointing at a model it cannot explain. The FCA has signalled that good and poor practice on explainability, audit trails and human oversight will be published later in 2026. Firms deploying AI without decision logs are building tomorrow's enforcement exhibits.

03

Data protection law has hardened underneath

The automated decision-making provisions of the Data (Use and Access) Act 2025 have been in force since February 2026, sitting alongside UK GDPR. Supervisory expectations stack on top of these obligations, not in place of them.

THE PRACTICAL READING

The regulator will not stop you experimenting — it is actively encouraging live testing. But a pilot with no data fence and no audit trail is not a pilot. It is an unmanaged exposure with a project name.

SECTION 04

The five properties of a real data fence.

Data fencing is often sold as a product. It is better understood as an architectural decision: no AI interaction happens except through a governed gateway the firm controls. Everything else follows from that single choice.

01 One front door

Every prompt, from every user and application, passes through a single gateway. Direct access to public endpoints is blocked at the network level. If there are two routes to a model, one of them is a leak.

02 Classification before transmission

Data is classified and screened at the gateway, before it leaves the boundary. Client identifiers and material nonpublic information are detected and redacted, tokenised or blocked automatically — not by asking the user to check.

03 Model allowlisting

Only approved models, on approved terms, with approved data handling, are reachable. Zero data retention and UK/EU residency stop being aspirations in a policy and become properties enforced in the pipe.

04 Full decision logging

Every prompt, response, model choice and routing decision is logged and retained. This turns AI usage from an unauditable behaviour into a supervisable activity — the audit trail the regulator increasingly expects.

05 Egress symmetry

Outputs are screened as well as inputs. A model fed sensitive context can reproduce it somewhere it should not go. The fence works in both directions or it is not a fence.

→ **Capability is not restricted**

Users still reach frontier models for the work that justifies them. What changes is that the firm — not the individual — decides what data travels, where it goes, and what record remains.

SECTION 05

Starve shadow AI — do not chase it.

Firms usually attack shadow AI with detection: scanning traffic, blocking domains, disciplining offenders. Detection is necessary but it is the smaller half of the answer. Shadow AI persists wherever the sanctioned route is slower, weaker or more annoying. Eliminate that gap and shadow usage collapses on its own.

MOVE ONE

Discover honestly

Map actual AI usage across the firm, including the embarrassing parts, before designing the sanctioned estate. The tools staff are sneaking are a free requirements document.

Map the real estate.

MOVE TWO

Make the sanctioned route the best route

Give users faster access to better models through the governed gateway than they could get on a personal account. When the compliant option is also the superior option, enforcement becomes almost incidental.

Best route, not just safe route.

MOVE THREE

Close the side doors

With a credible sanctioned route live, block direct access to public endpoints on corporate devices and networks. Blocking without an alternative breeds workarounds; blocking after a better alternative exists finishes the job.

Block after, never before.

THE SEQUENCE MATTERS

Detection finds the problem. A better sanctioned route removes the incentive. Blocking closes what remains — but only once the alternative is genuinely superior.

SECTION 06

The same gateway pays for itself.

There is a second reason to build the governed gateway now, and it has nothing to do with compliance. The commercial model underneath enterprise AI is shifting from subsidised per-seat licensing toward consumption-based, token-metered billing.

Firms budgeted for licences, they will be billed for usage

And usage is exploding: agentic workflows can consume many multiples of the tokens of a simple chat interaction. The FinOps Foundation's *State of FinOps 2026* found **73 per cent** of enterprises reported AI costs exceeding their original projections, and AI cost management is now the single most sought-after skill among technology-spend teams. Falling token prices have not helped, because volume is growing faster than prices are falling.

The dominant cause is architectural

Most estates default every query, however trivial, to the largest frontier model available. Summarising a paragraph does not need the same machinery as multistep reasoning over a regulatory filing, yet both are routinely priced identically because nothing in the estate distinguishes them.

WHERE GOVERNANCE EARNS ITS KEEP TWICE

The same control point that classifies data and logs decisions can also route intelligently — cutting token spend, in many cases by more than half, with no drop in output quality.

> **Governance and cost control are usually pitched as competing priorities. Route everything through one gateway and they become the same project.**
>
> — FORTAY CONNECT

SECTION 07

What good looks like: a ninety-day path.

None of this requires a transformation programme. The firms doing it well treat it as a focused engineering and governance exercise.

DAYS 1–30

Discovery & design

Map real AI usage, classify the data flows involved, and agree the model allowlist, data-handling terms and routing policy — with risk and compliance at the table from day one.

Discover, classify, agree policy.

DAYS 31–60

Build the fence

Stand up the gateway, wire in classification, redaction and logging, and connect the approved models. Migrate the highest-volume use case first to prove the experience is better, not just safer.

Build, wire, migrate.

DAYS 61–90

Eliminate & optimise

Close direct endpoint access, switch on intelligent routing, and hand risk a live dashboard: what is asked, what data was fenced, which models answered, and what it cost.

Close, route, report.

THE ESTATE IN ONE SENTENCE

Every interaction goes through one governed door, nothing sensitive leaves, everything is logged, and every query runs on the cheapest model capable of doing it well.

GET IN TOUCH

Risk, regulation & cost, built in from the start.

Fortay Connect is a UK-based CX and AI advisory practice working extensively with regulated financial services firms. We help organisations design, govern and scale AI across the client experience — from pilot to production — with risk, regulation and commercial sustainability built in from the start rather than retrofitted after the first incident.

If your firm is planning its next phase of AI projects, or has rolled out tooling and is now asking the governance and cost questions this paper raises, we would welcome a conversation — a short, no-obligation call on where you are, where the market is heading, and how to close the gap.

Book a 20-minute conversation → **Email Fortay Connect**

Mark Taylor · Fortay Connect
mark@fortayconnect.com · 0161 240 3411
Independent advisory for regulated contact centre organisations.