April 2026

Artificial Intelligence in Delegated Authority.

The market forces, regulatory imperatives, and strategic architecture that define AI in UK insurance.

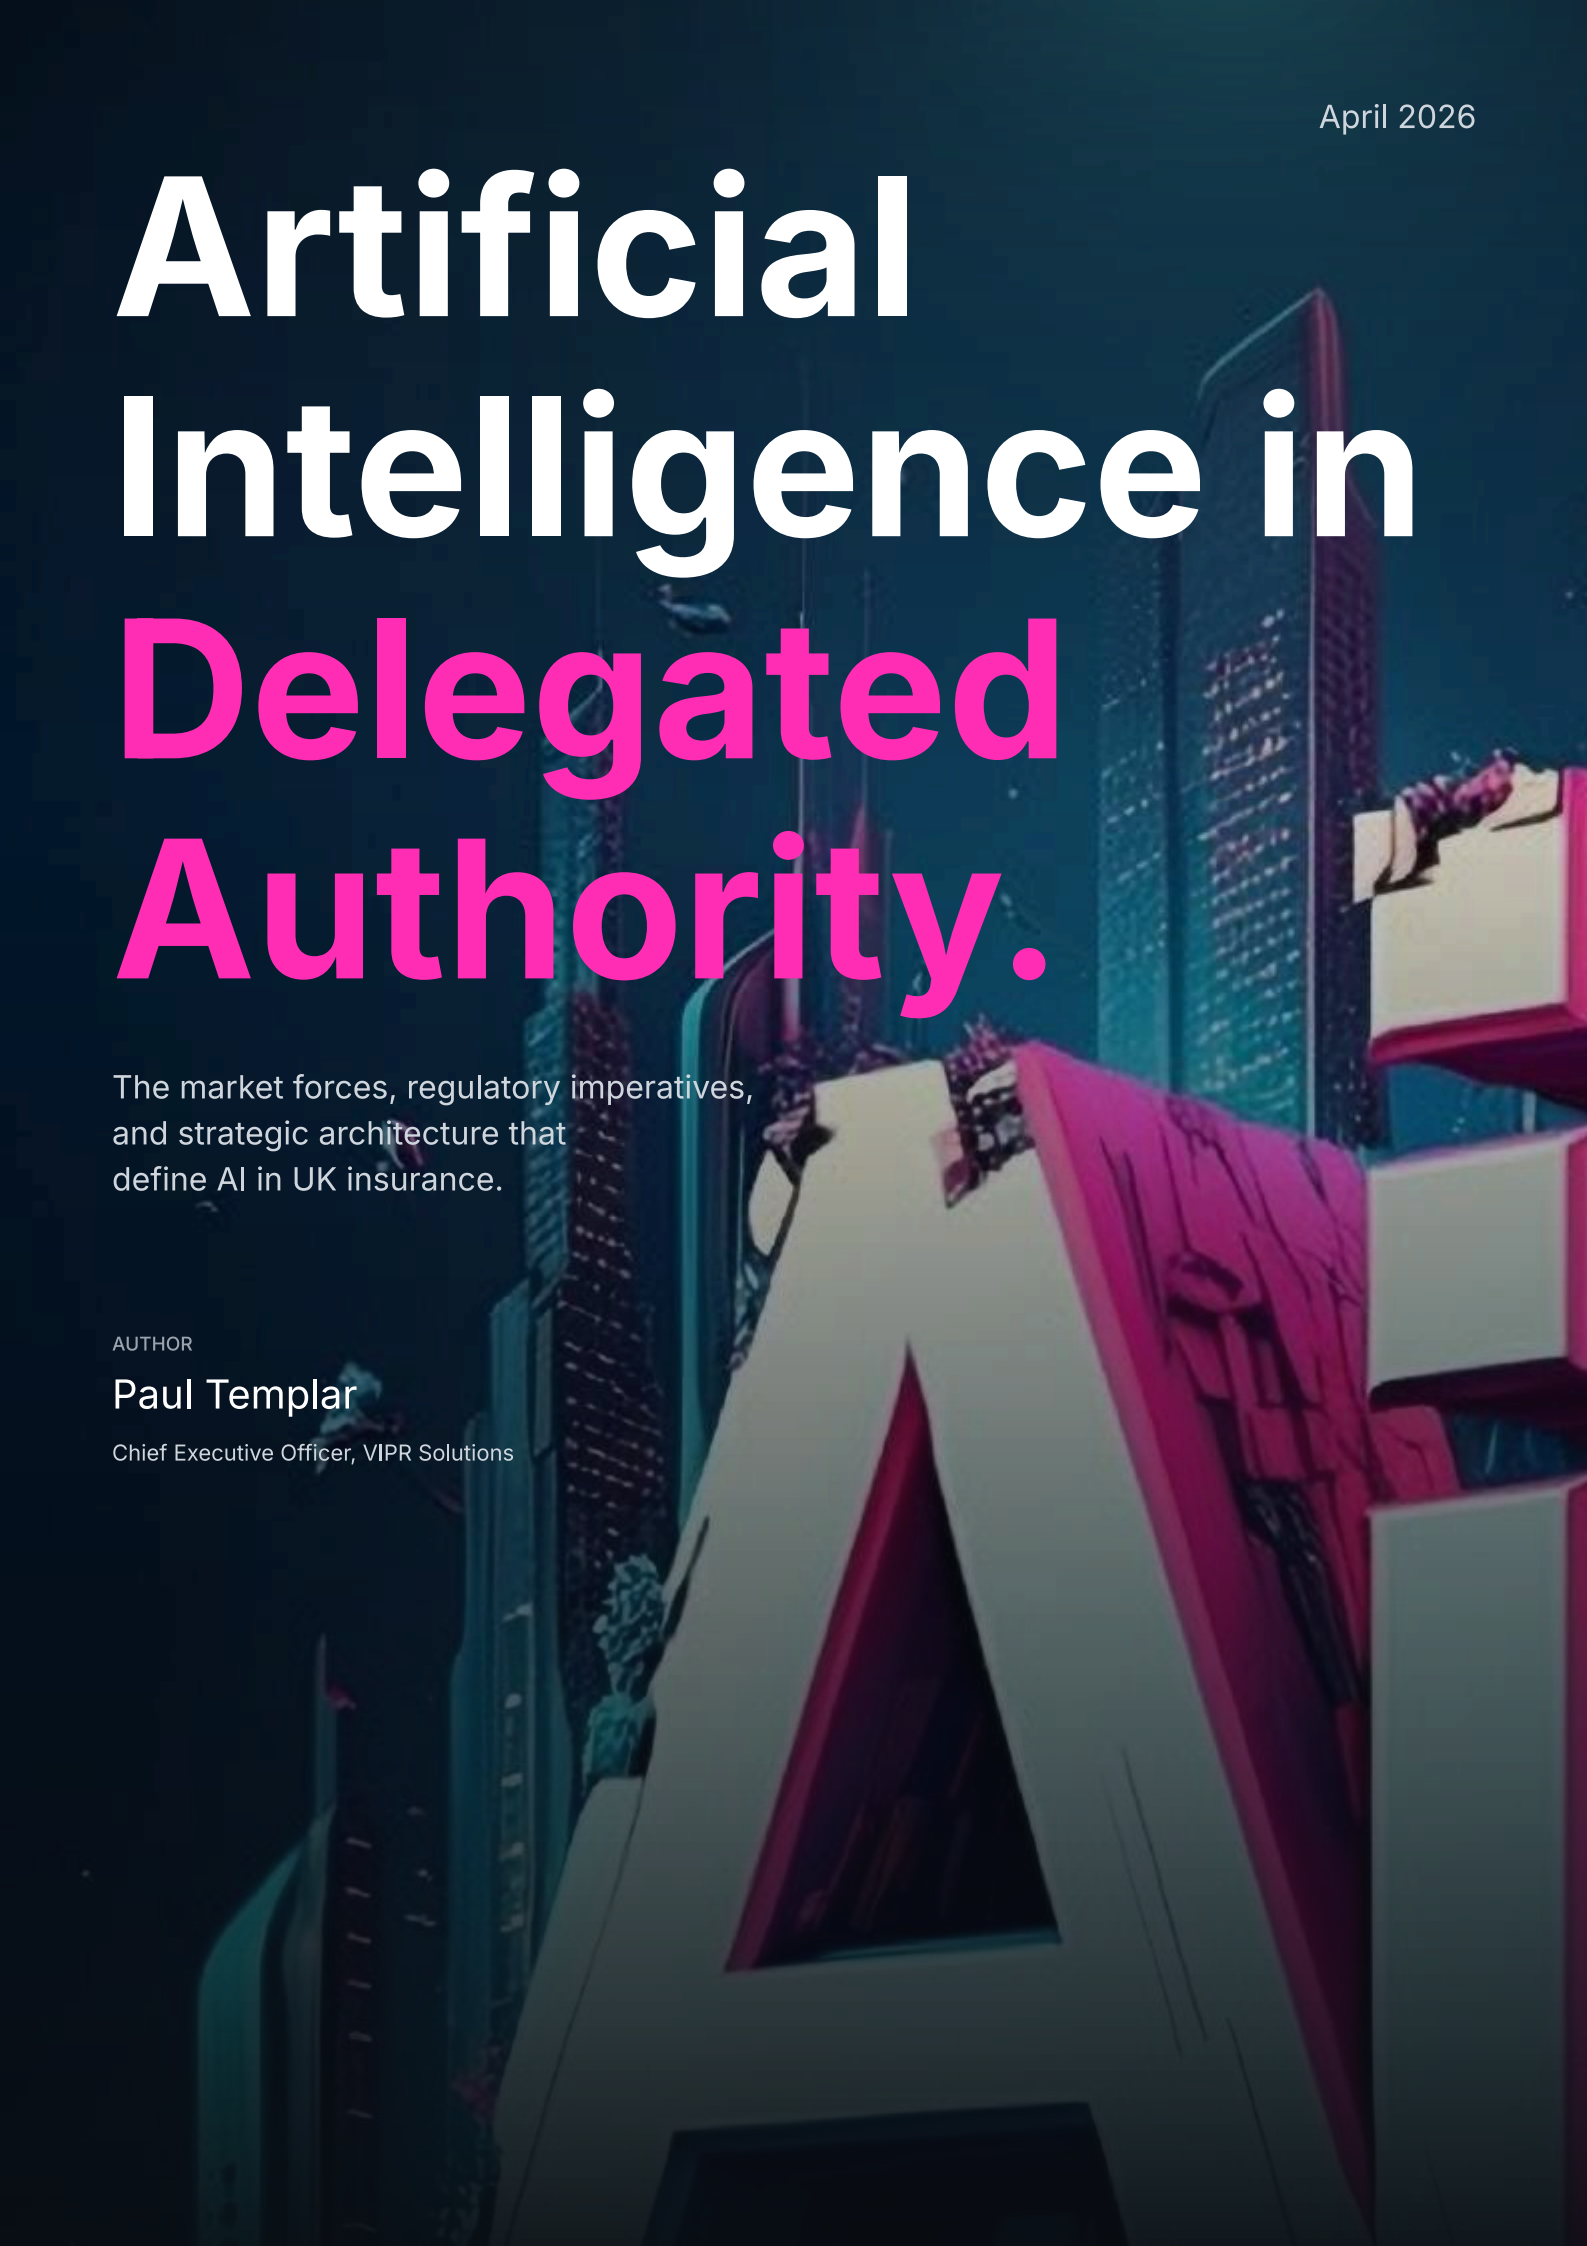

AUTHOR

Paul Templar

Chief Executive Officer, VIPR Solutions

ABOUT THE AUTHOR

PAUL TEMPLAR

Chief Executive Officer, VIPR Solutions

Paul Templar leads VIPR Solutions, the leading AI-enabled technology platform for delegated authority intelligence. Over £10 billion in gross written premium flows through VIPR annually, and more than 55% of Lloyd's Managing Agents run on the platform.

VIPR serves managing agents, coverholders, Lloyd's brokers, and DCAs across the London market and internationally, supported by offices in the UK and the United States.

"Specialist platforms with deep domain intelligence are more valuable in an AI-driven market, not less.

Inside this edition.

Three forces are converging.

Three forces are converging that will reshape the UK delegated authority market over the next 12 to 24 months. Each is significant on its own. Together, they define the strategic environment in which every managing agent, coverholder, Lloyd's broker, and DCA must operate.

01

AI is redefining which software businesses have real moats.

In February 2026, roughly one trillion dollars was wiped from software stock valuations in a period now known as the SaaSpocalypse. The median SaaS revenue multiple has dropped from above 7x to below 5x. The question every board is asking: can our technology platform be replicated by AI?

02

The regulator has drawn a direct line between DA oversight and AI.

The FCA's inaugural Regulatory Priorities: Insurance report, published 24 February 2026, simultaneously expands the review of delegated authority oversight and evaluates how firms use AI across underwriting, claims and consumer services. They converge on a single requirement: clean, structured, auditable data.

03

The architecture decision matters more than the AI itself.

In regulated financial services, the distinction between AI-native systems and Deterministic Process + AI systems is the difference between a product that can be audited and one that cannot. The companies that survive the SaaSpocalypse are those with moats that cannot be vibe-coded: regulatory compliance, proprietary data, multi-party workflows, and client trust built over years.

Three numbers that define the moment.

WIPED FROM SAAS VALUATIONS

$1tn

February 2026, the SaaSpocalypse. Median SaaS revenue multiple fell from above 7× to below 5×.

FCA PRIORITIES PUBLISHED

24 Feb

2026. The FCA's inaugural Regulatory Priorities: Insurance report, one document replacing more than 40 portfolio letters.

REGULATORY WINDOW

9-12mths

The window for managing agents, coverholders and DCAs to demonstrate audit-ready oversight before the DA review reports in early 2027.

Specialist platforms with deep domain intelligence, deterministic processing cores, and AI-native intelligence layers are more valuable in an AI-driven market, not less.

The economics confirm it. The regulatory environment demands it. And the competitive environment rewards it. Intended for the C-suite executives and board members of managing agents, coverholders, Lloyd's brokers, and DCAs evaluating their technology investments, regulatory readiness, and AI strategy.

NOV 2022 → FEB 2026

36 months that changed everything.

To understand where AI stands today in insurance, and where it is heading, it helps to understand the speed at which the market has shifted. What follows is not a technology history lesson. It is the context for a strategic decision every C-suite executive in the UK delegated authority market will make in the next twelve months.

Each generation of AI has expanded not just what the technology can do, but how independently it can do it.

NOV 2022	**ChatGPT launched.** 1M users in 5 days, 100M in 2 months. Every boardroom conversation about technology investment changed.
MAR 2023	**GPT-4 and Claude arrive.** The frontier becomes a two-horse race. Multi-step business reasoning is now commercially viable.
2023, 24	**The pilot era.** Conning 2025: 90% of insurers evaluating GenAI; only 42% had deployed. The gap was always data.
MAR 2024	**Claude 3.** Longer context windows mean entire policy documents and bordereaux files in a single pass.
LATE 2024	**The reasoning breakthrough.** OpenAI o1 and Claude extended thinking. Reserve calculations, regulatory analysis, exception triage become tractable.
FEB 2025	**Claude Code.** AI as engineering colleague. Product development cycles accelerate by an order of magnitude.
EARLY 2026	**The agentic era.** AI moves from a tool you ask questions to, to an agent that takes actions on your behalf.
FEB 2026	**The SaaSpocalypse.** Roughly $1 trillion wiped from software valuations. AI is now the product, or the threat.

HOW AI ADOPTION EVOLVED

UK insurance has moved through four distinct phases.

The pace of change is accelerating, not slowing. The question for every C-suite executive is not whether to engage with AI, but whether their data infrastructure is ready for what comes next.

PHASE 01

25%

20 23

Document processing

Extracting data from unstructured documents. Policy intake, claims forms, bordereaux parsing. Useful, but narrow, most implementations were glorified OCR with a language model attached.

PHASE 02

50%

20 24

Workflow automation

Routing claims, pre-populating submissions, flagging anomalies. Travelers ($1.5bn annual technology spend), Liberty Mutual (16 GenAI use cases) and Nationwide ($100m/year on AI) showed real results. Most of the UK market remained in pilot.

PHASE 03

75%

20 25

Decision support

Reasoning models analyse risk scenarios, identify portfolio concentrations, recommend pricing adjustments. BCG reports 36% underwriting efficiency gains from multi-agent AI deployments.

PHASE 04

100%

20 26 → CURRENT FRONTIER

Supervised autonomy

AI manages entire workflows end-to-end: bordereaux validation, exception triage, regulatory reporting, FCA evidence packs, with humans providing oversight at key decision points.

PART ONE

The Market Context.

Why the SaaSpocalypse exposes the difference between shallow moats and deep ones, and why regulated vertical software survives.

The SaaSpocalypse and what it means for UK insurance technology.

The February 2026 sell-off described in the Executive Summary was not a momentary correction. It exposed a structural divide in the software industry: businesses with deep moats versus those without.

Investors are now applying a simple test to every software business in their portfolio: can this be vibe-coded? If the answer is yes, the moat is thin regardless of how prominently AI features in the narrative.

The numbers tell the story. The median SaaS revenue multiple has dropped from above 7x to below 5x. SaaS applications per company are down 18% in 2025. And Gartner predicts that more than 40% of agentic AI projects will be cancelled by end of 2027.

The conversation has flipped. The question on every board is no longer "how do we add AI?", it is **"what protects us from AI commoditisation?"**

$1tn+
Wiped from software valuations
February 2026

<5x
Median SaaS revenue multiple
down from 7x+

−18%
SaaS apps per company
2025 decline

40%+
Agentic AI projects cancelled
by end 2027 (Gartner)

95%
GenAI pilots fail to deliver
P&L impact (MIT)

CALLOUT · WHAT VIBE CODING CAN & CANNOT DO

Vibe coding is effective at generating simple, self-contained applications: internal tools, dashboards, CRUD interfaces, basic automation. It is **not effective** at generating software that must comply with financial services regulation, integrate with decades-old data infrastructure, handle multi-party workflows with contractual obligations, or maintain auditable decision trails across millions of transactions.

PULL QUOTE

"The SaaSpocalypse is real for horizontal SaaS with shallow moats. It is not real for domain-specific software in regulated industries.

WH Y REGULATED VERTICAL SOFTWARE SURVIVES

The companies that survive the SaaSpocalypse are those whose moats cannot be vibe-coded. Four moats hold up.

01

Regulatory compliance

Years of accumulated regulatory knowledge embedded in software logic, EU AI Act, Solvency II, IDD, DORA, FCA Consumer Duty, Lloyd's coverholder standards, and (for cross-border programmes) NAIC frameworks.

02

Proprietary data

Platforms at the centre of multi-party data flows accumulate proprietary data that cannot be replicated by a new entrant. AI performance compounds with that data.

03

Multi-party workflows

Networks of managing agents, coverholders, DCAs, brokers, reinsurers, and regulators with contractual obligations. Production systems spend most of their complexity budget on the 20% of cases that are not straightforward.

04

Client trust compounds

In a zero-tolerance accuracy environment, switching costs are real, not artificial. Trust is earned transaction by transaction, and it is the hardest moat for any competitor to overcome.

The moat is built from compliance, proprietary data, and operational complexity that cannot be prompted into existence.

Regulatory compliance cannot be vibe-coded. Meeting the EU AI Act, Solvency II, IDD, DORA, FCA Consumer Duty obligations, and (for cross-border programmes) NAIC AI Model Bulletin requirements demands years of accumulated regulatory knowledge embedded in software logic.

Proprietary data is a structural moat. Platforms at the centre of multi-party data flows accumulate proprietary data that cannot be replicated by a new entrant. This data improves AI performance over time, creating a compounding advantage.

Multi-party workflows resist simplification. UK delegated authority operates through networks of managing agents, coverholders, Lloyd's brokers, DCAs, reinsurers, and regulators with contractual obligations, regulatory requirements, and legacy integration points.

Client trust compounds. In a zero-tolerance accuracy environment, switching costs are real, not artificial. Trust is earned transaction by transaction, and it is the hardest moat for any competitor to overcome.

THE TAKEAWAY

AI does not disrupt platforms that have distribution, domain expertise, and encoded intelligence. It disrupts the manual processes those platforms replace.

The Regulatory Imperative.

The FCA has put DA oversight and AI on the same page. The 9, 12 month window has begun.

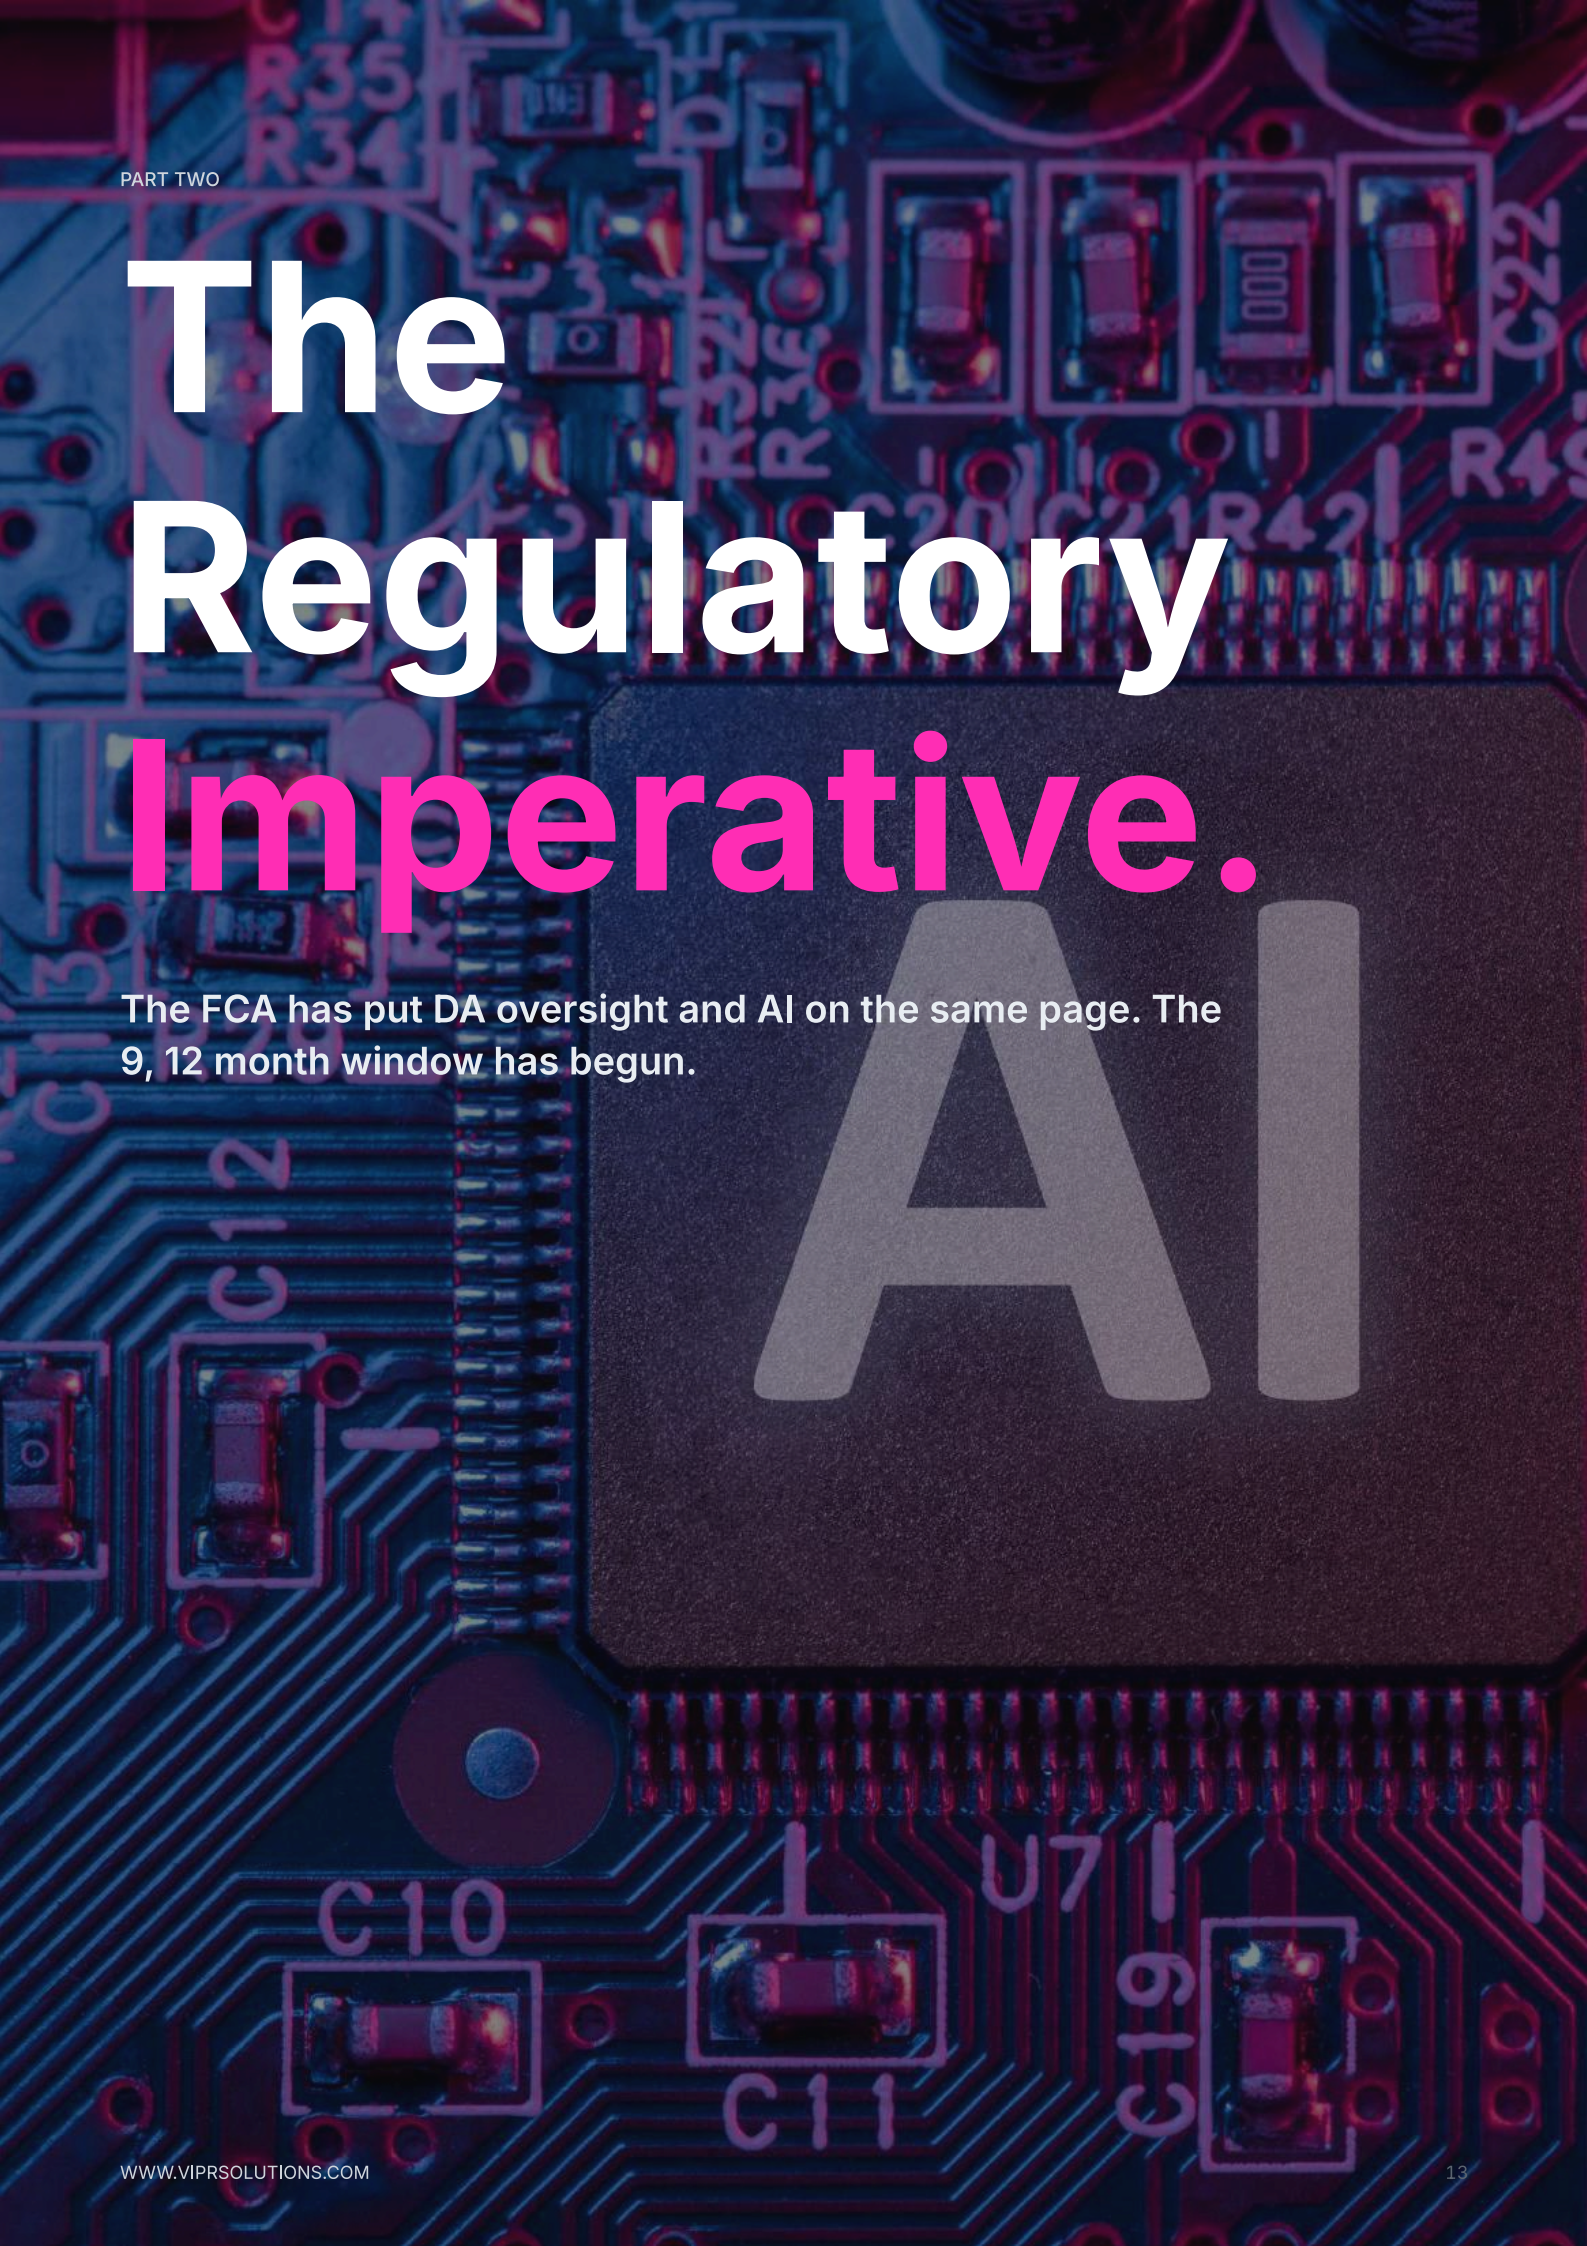

The FCA has put DA oversight and AI on the same page.

On 24 February 2026, the Financial Conduct Authority published its inaugural Regulatory Priorities: Insurance report, a single document replacing more than 40 portfolio letters that previously set the agenda for different parts of the insurance market. Two priorities deserve the attention of every firm operating in delegated authority.

01 · DA OVERSIGHT REVIEW

The FCA is expanding its review of DA oversight.

Examines firms' oversight of outsourced claims processes across different delegated authority models and remuneration arrangements. **Commences Q2 2026, with findings expected early 2027.**

02 · AI LIVE TESTING

The FCA is evaluating how firms use AI.

Across underwriting, claims and consumer services. **An evaluation report from the AI Live Testing programme is expected by end of 2026.**

WHERE THEY CONVERGE

These are not separate initiatives. They converge on a single question: does your organisation have the data infrastructure to demonstrate **effective oversight of your delegated authority book**and to **adopt AI responsibly?**

DA OVERSIGHT · TH E PRINCIPAL CANNOT DELEGATE ACCOUNTABILITY

Delegated authority premiums doubled from £10.4 billion in 2018 to £22.1 billion in 2023, and DA is projected to exceed 50% of total London market premium by 2027. The FCA's decision to dedicate specific supervisory attention to DA oversight reflects the scale and complexity the market has reached.

Capacity providers remain accountable for everything written under their authority.

DA PREMIUM 2018

£10.4bn

Delegated authority premium written across the London market.

DA PREMIUM 2023

£22.1bn

More than doubled in five years, driving the FCA's decision to focus.

PROJECTED 2027

>50%

DA share of total London market premium by 2027.

BOARD EXPECTATIONS

Boards will need to evidence robust oversight frameworks, strong management information, and clear lines of responsibility under the Senior Managers and Certification Regime.

CONSUMER OUTCOMES

Customer outcomes under DA must equal in-house provision. The FCA is not accepting lower standards because claims handling sits outside the principal.

REMUNERATION

Remuneration structures are under the microscope. The review will examine whether incentives within DA arrangements align with fair claims outcomes.

GOVERNANCE PRECEDES DEPLOYMENT

AM Best identified AI as the number one emerging risk for DUAEs, noting that many have not yet implemented formalised governance around their AI model inventory.

THE DOOR IS OPEN · EVIDENCE REQUIRED

The FCA is not anti-AI. It requires evidence.

The regulator has built a full AI innovation infrastructure, the Supercharged Sandbox (in partnership with NVIDIA), AI Live Testing for firms ready to deploy in live markets, and the AI Spotlight publishing real-world insights into firm experimentation.

> **"The door is open, firms just have to walk through it.**
>
> Sarah Pritchard, FCA Executive Director of Markets & International

Yet only 208 of nearly 3,000 Sandbox applicants were insurance-related, despite insurance firms enjoying a 50% acceptance rate compared to 31% overall.

I. CONSUMER DUTY ALIGNMENT

Non-negotiable. Any AI system touching customer journeys must demonstrate good consumer outcomes with governance, explainability, and outcome testing.

II. BIAS MONITORING MANDATORY

AI producing unintended exclusion or discriminatory outcomes will be treated as a Consumer Duty breach.

III. GOVERNANCE PRECEDES

Document your AI model inventory and governance framework before deployment, not after.

TH E 9, 12 MONTH WINDOW

Firms that act now will be ahead of the curve when the findings are published. Firms that wait will be reacting to regulatory findings rather than shaping their response.

Q2 20 26	Q4 20 26	Q1 20 27
DA oversight review commences. EU AI Act takes effect (Aug 2026) for high-risk insurance use cases.	FCA AI Live Testing evaluation report. Firm experimentation outcomes published.	DA oversight review findings published. AM Best DUAE outlook revisit.

PART THREE

The Architecture Decision.

AI-Native vs Deterministic Process + AI is not a technical preference. It is the difference between a product that can be audited and one that cannot.

AI-Native vs Deterministic Process + AI.

These are not points on a spectrum. They are fundamentally different architectural decisions with different implications for reliability, auditability, cost, and risk.

"Remove the AI from an AI-native system and the system does not function. Remove the AI from a Deterministic Process + AI system and the deterministic core continues to operate.

	OPTION A · AI-NATIVE	OPTION B · DETERMINISTIC PROCESS + AI
Core execution	AI is in the critical path. Probabilistic models as primary decision layer.	Deterministic logic handles core execution. AI enhances specific functions.
Output behaviour	Same input can produce different outputs. Confidence scores and fallback chains are structural.	Same input always produces the same output. Reproducibility and auditability built in.
Failure mode	Probabilistic: hallucination, drift, adversarial manipulation. Harder to detect.	Deterministic: bugs, logic errors, data issues. Traceable and reproducible.
Audit trail	Requires purpose-built observability: confidence scoring, decision logging.	Inherent in the deterministic logic. Every decision traceable to a rule.
Regulatory fit	**Difficult.** "The model decided" is not acceptable for the FCA.	**Strong.** Decision logic is explainable, reproducible, and documentable.

Why Deterministic Process + AI is the right architecture today.

THE PATTERN THAT WINS

AI proposes, rules constrain, humans approve, and a complete audit trail records every step.

TH E REGULATORY REALITY

Regulated firms must explain why a decision was made, reproduce it, and prove it was not discriminatory. The EU AI Act (effective August 2026) classifies AI systems used for risk assessment in insurance as **high-risk**requiring documented decision paths, human oversight, and bias mitigation. Deterministic systems meet these requirements by default.

TH E RELIABILITY GA P

Gartner predicts over 40% of agentic AI projects will be cancelled by end of 2027. A separate MIT-linked study found 95% of enterprise generative AI pilots fail to deliver measurable P&L impact. These are arguments against AI as the primary execution layer in systems where accuracy is non-negotiable.

TH E DATA PROBLEM

Most high-value data in UK insurance lives on legacy systems never designed to serve as AI training pipelines. A Deterministic Process + AI architecture **tolerates messy inputs** because the deterministic pipeline validates, normalises and structures data before any AI layer touches it.

TH E HYBRID ARCHITECTURE THAT WORKS

The binary framing misses the architecture that is actually winning in regulated financial services: a **deterministic core with an AI-native intelligence layer on top.**

Data ingestion & normalisation LAYER 01	DETERMINISTIC	Bordereaux validation, format conversion, schema mapping
Core processing LAYER 02	DETERMINISTIC	Premium calc, reserve computation, regulatory reporting
Exception handling LAYER 03	AI -ASSISTED	Flagging anomalies, suggesting corrections, routing exceptions
Intelligence & insight LAYER 04	AI -N ATIVE	Natural language querying, anomaly detection, portfolio analytics

Each phase proves itself before the next begins.

The case against AI-native data processing is a timing argument, not a permanent one. VIPR is already executing a deliberate transition, deployed in layers.

IN PRODUCTION

01

Schema mapping

AI maps incoming data to target schemas, resolves ambiguities. Shadow mode validation against deterministic rules.

20 25, 26

UNDERWAY

02

Exception handling

AI identifies and routes exceptions requiring manual intervention. Confidence scoring and human-in-the-loop review.

20 26, 27

PLANNED

03

Validation logic

AI applies validation rules with human oversight on exceptions. Auditable reasoning chains, regulatory acceptance.

20 27, 29

FUTURE

04

Regulated outputs

AI generates regulatory submissions, audit reports. Full explainability, regulatory safe harbour.

20 29 +

Each phase depends on the one before it. VIPR's clients gain AI capabilities as they mature, without taking on the risk of premature automation in regulated processes. Skipping phases introduces risk that regulators, clients, and E&O insurers will not tolerate.

PART FOUR

VIPR's Position.

Why AI makes a specialist platform with deep domain intelligence more valuable, not less.

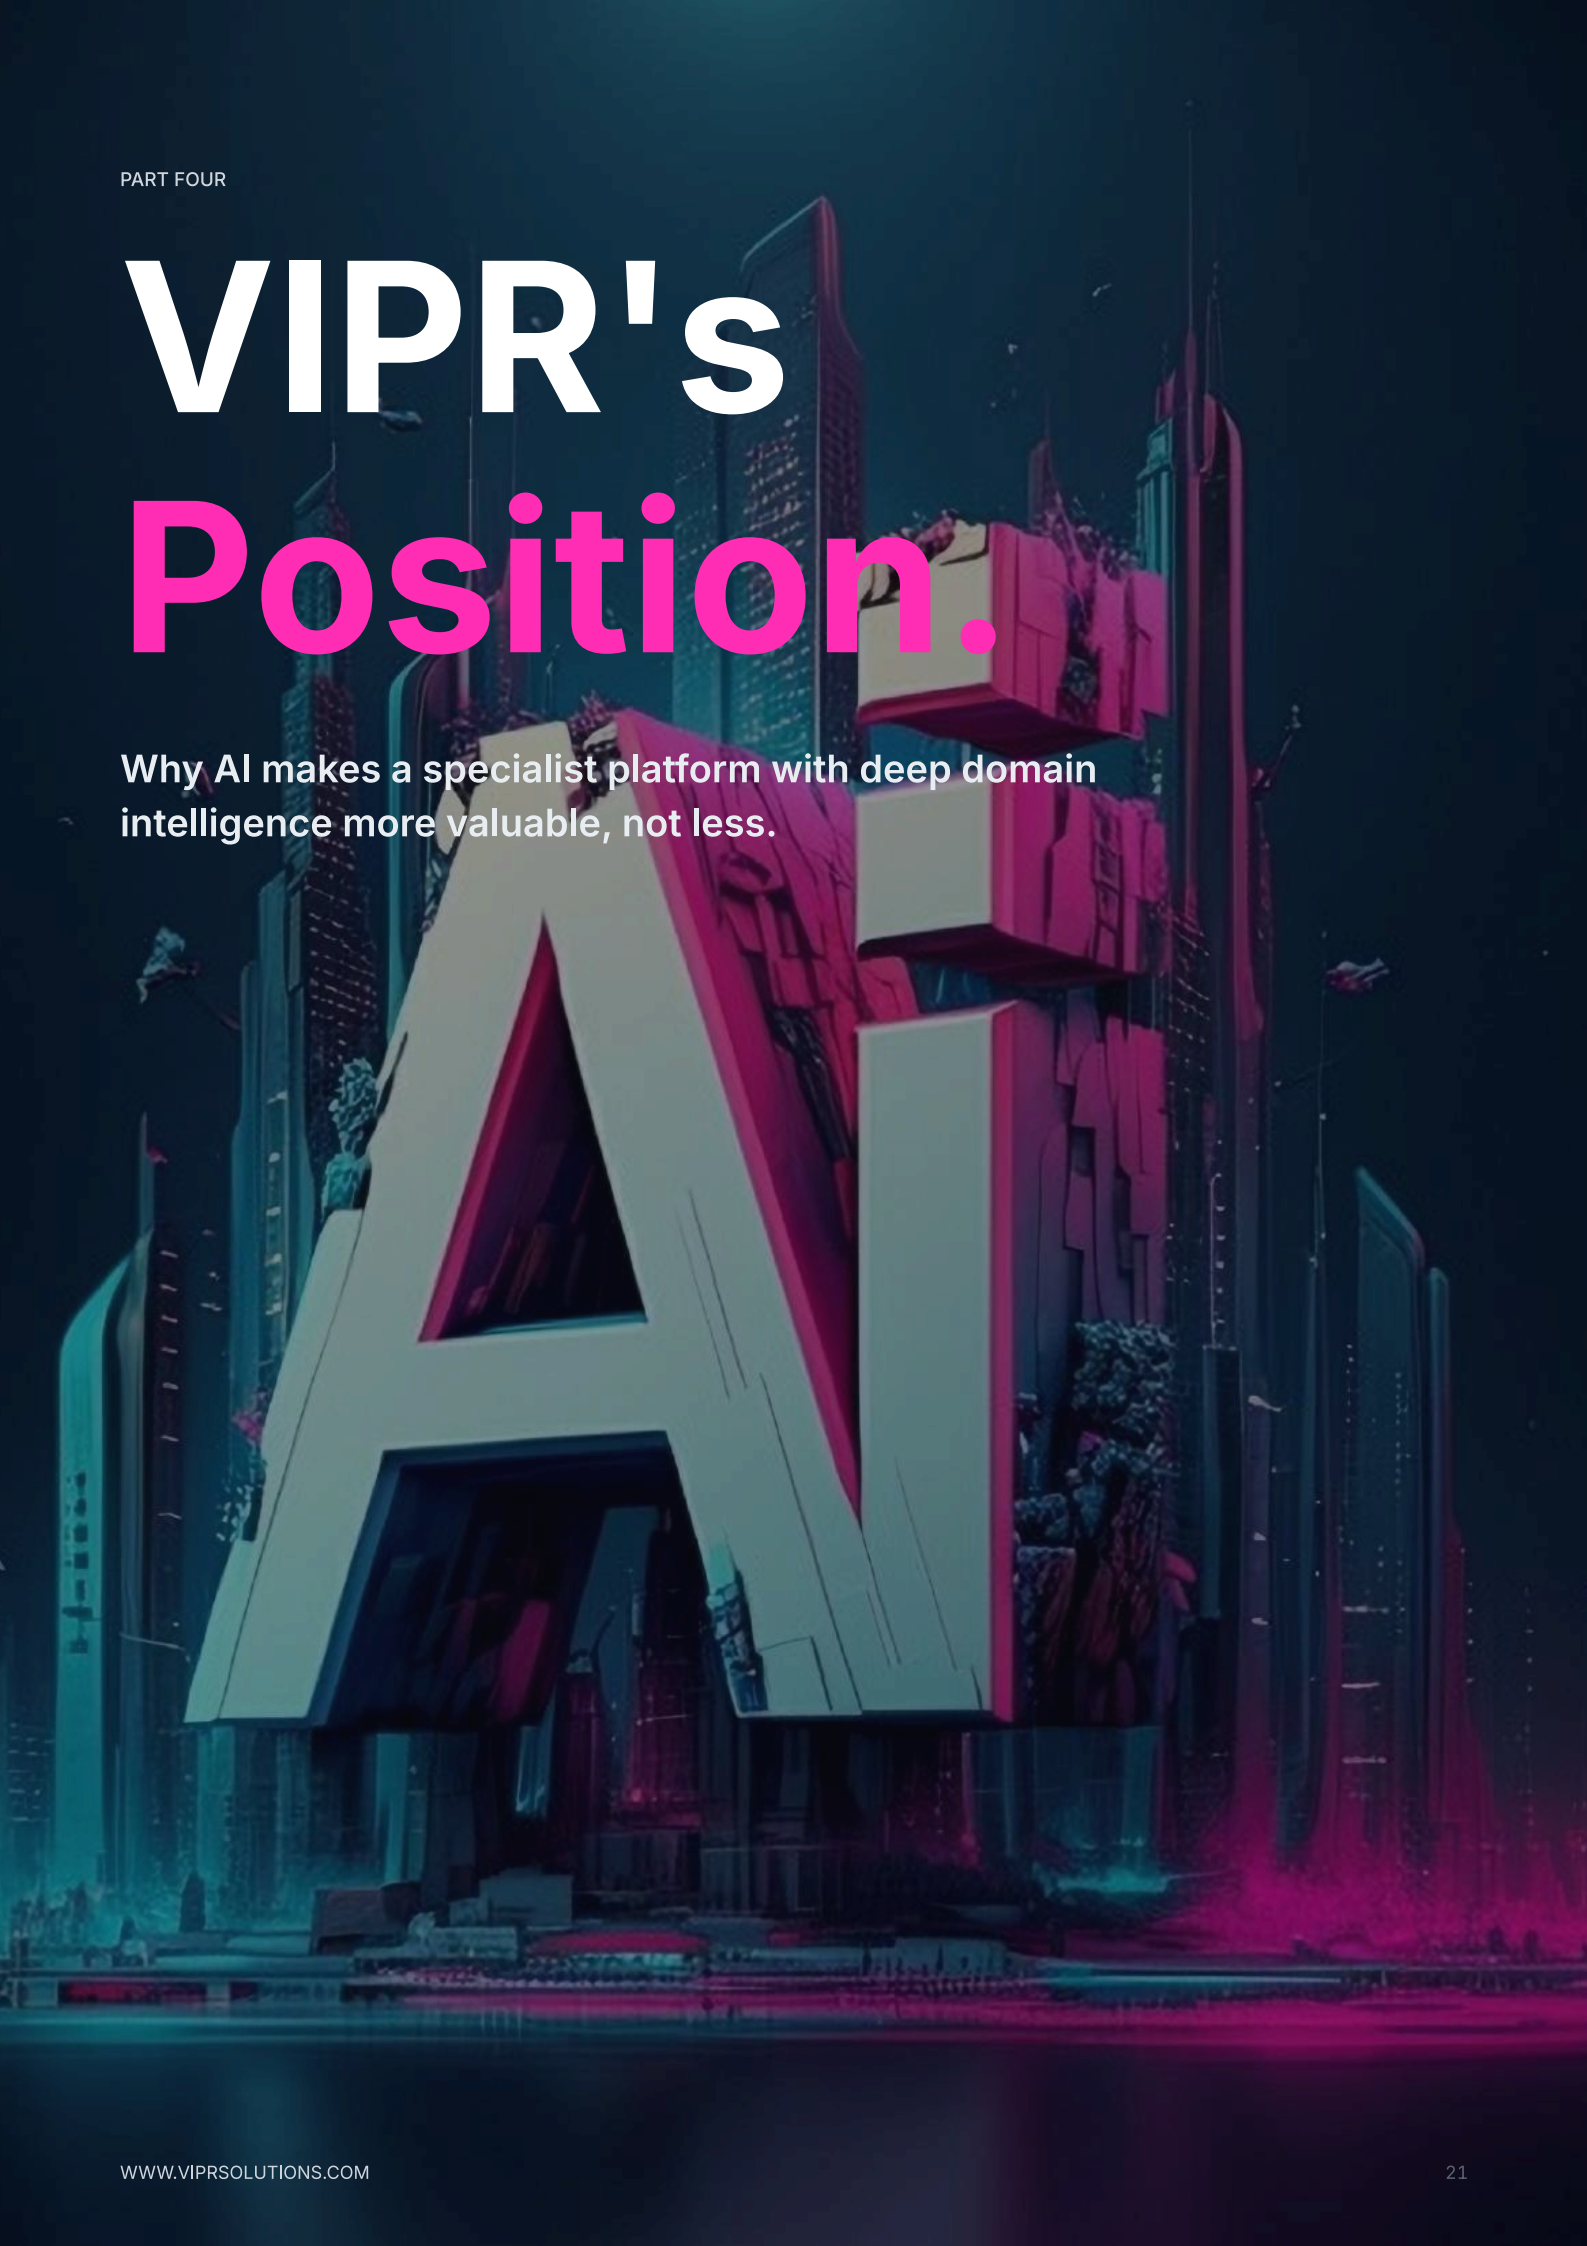

What it would cost to replace VIPR with AI alone.

We modelled what it would cost a single large UK customer to process delegated authority data using AI alone: 1,000 bordereaux files per month, averaging 5,000 rows and 40 columns each. Schema mapping, transformation, lookups, FX and tax, validations, consistency checks, historical comparison.

APPROACH	ANNUAL COST*	VIABLE?
Pure AI No specialist platform	**~£1m+ p.a.**	NOT VIABLE
In-house build Custom platform	**£500k–£1m p.a.**	YEARS TO DELIVER
VIPR, platform + AI Live today, market-trusted	**£200–300k p.a.**	OPERATIONAL

*Per UK customer, processing 1,000 bordereaux files per month. Both pure-AI and hybrid approaches assume perfect accuracy, which no AI model delivers.

10–20 MILLION TOKENS PER FILE

Cost is not even the binding constraint. A routine bordereaux file with 40,000 rows translates to 10, 20 million tokens, far exceeding the largest AI context window available today. Some client files exceed one million rows.

Even if model costs fell by 90%, the pure-AI approach would still cost over £100,000 per customer per yearand would still fail on file sizes VIPR handles every day.

Rather than replacing the platformAI extends it.

Today, VIPR delivers fully automated, zero-touch handling for clean bordereaux data. **Roughly 70% of transactions** flow through without human intervention. AI is now being deployed against the remaining 30% of exceptions and edge cases.

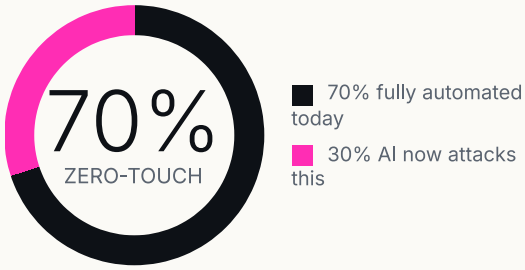

Intelligent schema matchingnew schemes onboarded in minutes, not days.

Predictive exception handlingflags anomalies before they become errors.

Automated validationcatches issues before they propagate downstream.

FIVE REINFORCING ADVANTAGES · WHY VIPR COMPOUNDS

01

Encoded domain intelligence

Over 4,000 schema mappings and 17 years of scheme-specific configuration & validation logic. Cannot be scraped or replicated.

THE DATA MOAT

02

Market leadership

Over 55% of Lloyd's Managing Agents run on VIPR, alongside large UK and US carriers, brokers and reinsurers.

THE PROOF

03

Domain complexity as a moat

Bordereaux is a regulatory, contractual and operational problem. 10,000+ delegated authority contracts processed monthly.

THE BARRIER

04

Switching costs & trust

Years of historical data, retrained teams, rebuilt integrations, re-validated compliance, all live with VIPR today.

THE LOCK-IN

05

VIPR's AI is accelerating

Built on real operational context, not general-purpose models. Learns faster, hallucinates less, with expert humans in the loop.

THE COMPOUNDING EDGE

"AI does not disrupt platforms that have distribution, domain expertise, and encoded intelligence. It disrupts the manual processes those platforms replace.

VIPR is on the right side of that equation.

PART FIVE

What this means for you.

For VIPR clients already running on the platform, and for those evaluating their next move.

Two audiencesone position.

IF YOU ARE A VIPR CLIENT

AI capabilities are designed to enhance the platform you use today.

No new system to buy or integrate. Embedded in workflows your teams already use.

Faster onboarding of new schemes. Days of manual mapping become minutes of intelligent matching.

Fewer exceptions, resolved faster. AI learns from historical patterns and routes complex cases.

Higher data quality without additional effort. Predictive validation catches issues before they propagate.

No risk to existing workflows. AI operates alongside, not in place of, deterministic processing.

Costs improve as the AI matures. Gains flow directly to you through the platform.

IF YOU ARE EVALUATING

The case for a specialist platform has never been stronger.

Regulatory readiness. A specialist platform delivers clean, structured, auditable data by default, exactly what the FCA's DA review demands.

Proven economics. Augmenting a specialist platform with AI is a fraction of either alternative.

Competitive advantage. Better data infrastructure means better underwriting, pricing and operational decisions.

Time to value. 4,000+ schema mappings and 17 years of encoded intelligence, operational from day one.

The 9, 12 month regulatory window means firms need infrastructure that works nownot in 2028.

What can AI do, what should AI do, and who is accountable?

In a year when 78% of business leaders report that AI adoption is outpacing their ability to manage risk, this is not abstract. The FCA's 2026 priorities make it explicit: AI adoption is welcome, but only with evidence of good outcomes, explainability, and human oversight. VIPR's approach is built on three principles.

01

Purposeful automation.

AI is applied where it adds the most value: reducing manual effort on the 30% of exceptions that currently require human judgement. The deterministic platform continues to handle volume processing with the accuracy and auditability you expect.

02

Full auditability and human oversight.

Every AI-assisted decision can be traced, reviewed, and explained. In a regulated environment, this is the minimum acceptable standard.

03

Measurable outcomes.

Every AI capability is measured against operational KPIs: processing time, exception rates, error reduction, and cost per transaction. If it cannot demonstrably improve operations, it does not ship.

Three forcesone direction of travel.

This document has traced a single thread through three forces reshaping the UK delegated authority market simultaneously. AI has moved from novelty to infrastructure in 36 months. The software market has repriced around that reality, punishing shallow moats and rewarding depth. The FCA has drawn a direct line between DA oversight and AI governance, giving the market a 9, 12 month window to demonstrate readiness. And the architecture question, AI-native versus Deterministic Process + AI, has moved from a technical preference to a strategic choice with regulatory, commercial, and operational consequences.

For managing agents, coverholders, Lloyd's brokers, and DCAs operating in the UK delegated authority space, these forces all point in the same direction. The firms that invest in AI-ready data infrastructure now will be positioned for both the regulatory environment and the competitive opportunity. The firms that treat data as an operational afterthought will find themselves on the wrong side of a market that has moved on without them.

VIPR's position is clear. The platform becomes more valuable as AI matures, not less. The domain intelligence encoded in seventeen years of operation is the foundation that makes AI effective in delegated authority. Without it, AI is expensive, inaccurate, and limited by technical constraints. With it, AI delivers faster processing, smarter exception handling, and lower operational costs, all through the platform your teams already know.

THE TAKEAWAY

AI does not disrupt platforms that have distribution, domain expertise, and encoded intelligence. It disrupts the manual processes those platforms replace.

Six imperatives. The next 12 months.

01

Map your data flows.

Understand where data enters, how it moves, where gaps exist. The FCA audit trail starts here.

02

Automate validation upstream.

Errors caught at point of entry are operational improvements; errors caught by the regulator are governance failures.

03

Build audit trails now.

Real-time validation with full traceability is a stronger answer than quarterly reviews.

04

Prepare AI governance frameworks.

Document your AI governance before the FCA's evaluation sets expectations.

05

Connect premium, claims and financial data.

Isolated data sets will not satisfy the regulator's interest in consumer outcomes.

06

Evaluate your architecture.

Determine whether your processing is auditable and your intelligence layer is positioned to differentiate.

ABOUT VIPR SOLUTIONS

VIPR Solutions is the global leader in delegated authority and bordereaux data management technology. Over £10 billion in gross written premium flows through VIPR annually, serving managing agents, coverholders, Lloyd's brokers, and DCAs across the London market and internationally. More than 55% of Lloyd's Managing Agents run on the VIPR platform.

2009	PE	SOC 2	UK·US
FOUNDED	BACKED · TENZING	CERTIFIED	OFFICES

THE CLOCK IS TICKING

Get your DA data readiness assessment before the Q2 review window closes.

AUTHOR

Paul Templar

Chief Executive Officer, VIPR Solutions

MORE

viprsolutions.com

paul@vipr.solutions