

DIGITAL EDGE

[illegible]

Exclusive Interview:

WILLIAM FOLEY

Stratio

The Trust Paradox:

Why Enterprise AI Fails Before It Starts





Behind the Solutions

Powering Data-Driven Transformation Across Industry



A conversation with

William Foley



The Trust Paradox:

Why Enterprise AI Fails Before It Starts

William Foley has watched brilliant AI models die in risk committees. Not because they failed, but because no one could defend them.

In the gleaming innovation labs of Europe's largest banks, data scientists are building AI systems that can assess creditworthiness in milliseconds, detect fraud with uncanny precision, and automate compliance workflows that once consumed entire departments. The models work. The accuracy is there. The business case is compelling. And then, somewhere between proof of concept and production deployment, these projects enter what Foley calls "governance purgatory": a months-long limbo where risk

committees, compliance officers, and regulators ask a single, devastating question: Can you prove this decision is defensible? More often than not, the answer is no. "We don't have an AI scaling problem," Foley tells me from Stratio's offices, his tone carrying the weight of two decades spent in the trenches of financial services transformation. "We have a trust scaling problem." It's a counterintuitive thesis in an industry obsessed with compute power, model sophistication, and algorithmic breakthroughs. But Foley's insight



Digital Edge's regular deep-dive into the major tech vendors and solution providers driving enterprise innovation.

Each feature uncovers the story behind the technology and services — from the challenges they solve to the teams behind the scenes and the customers they empower.

cuts deeper: the real bottleneck in enterprise AI isn't capability. It's the capacity to make decisions auditable, explainable, and defensible at the speed and scale that AI demands. In financial services, where AI decisions carry direct human consequences (credit denials, fraud flags, insurance claim rejections), regulators are rightfully demanding that institutions demonstrate not just that their models work, but that they can account for why they work, how they arrived at specific decisions, and what happens when those decisions are challenged three years down the line.

The institutions that crack this problem won't just deploy AI faster. They'll fundamentally reshape competitive advantage in financial services. The ones that don't will find their innovation labs producing expensive science projects that never leave the building.

FROM AUDIT TRAILS TO AI TRAILS

William Foley's career reads like a masterclass in understanding why governance matters before it becomes a crisis. He started at KPMG, learning the foundational discipline of audit: evidence, traceability, defensibility. "In audit, you learn very quickly that good intentions don't matter," he says. "What matters is whether you can reconstruct a decision, demonstrate the evidence that supported it, and defend it under scrutiny. That mindset, that rigour, is almost more relevant now than it ever was."

From KPMG, Foley moved to Santander, where he spent twelve years in progressively senior roles driving operational efficiency and, crucially, managing the nightmarish complexity of post-acquisition integration. Anyone who's worked in financial services M&A knows the pattern: two banks merge, and suddenly you're trying to reconcile customer data across incompatible systems, each with its own definitions, its own data quality standards, its own governance frameworks. "You very quickly learn the true cost of data fragmentation," Foley recalls. "It's not just technical debt. It's strategic paralysis. You can't make decisions with confidence because you don't trust the data underneath them."

It was during his tenure as Chief Data Officer of Santander's Innovation division that Foley validated what would become his central thesis: governance isn't a constraint on innovation; it's a precondition for it. "We built customer data foundations that became the platform for everything else," he explains. "Once you have trusted data, once you have governance embedded at the point of consumption rather than bolted on afterwards, suddenly you can move fast. You can experiment. You can scale. But you have to build that foundation first."

That experience, watching governance transform from bureaucratic overhead into strategic enabler, shapes everything Foley now brings to his role as Chief Product Officer at Stratio, where he's helping financial institutions build the infrastructure that makes AI not just possible, but defensible.



THREE SEISMIC SHIFTS

Foley identifies three fundamental transformations that have reshaped financial services over the past two decades, each one amplifying the importance of the others.

First, data moved from back-office cost centre to board-level strategic asset. "Twenty years ago, data was something you stored because you had to," Foley says. "Regulatory requirements, audit trails, historical records. It was overhead." Today, data is the raw material of competitive advantage: the foundation for personalisation, risk modelling, operational efficiency, and increasingly, autonomous decision-making. But that elevation brings scrutiny. When data becomes strategic, data quality becomes existential.

Second, regulatory compliance shifted from constraint to precondition for innovation. The wave of post-2008 regulation (Basel III, GDPR, MiFID II, and now the EU AI Act and Digital Operational Resilience Act (DORA)) hasn't just increased compliance burden. It's fundamentally changed the relationship between innovation and regulation. "You can't innovate your way around compliance anymore," Foley observes. "The regulators are too sophisticated, the penalties too severe, and the reputational risk too high. Compliance isn't the thing that slows you down. It's the thing that determines whether you can scale at all."

Third, and most profoundly, AI has shifted from analytical tool to autonomous decision-maker. "For decades, AI and machine learning were tools that helped humans make better decisions," Foley explains. "You'd run a model, get a recommendation, and a human would



review it and make the call. That's changing. We're moving toward AI agents that make decisions autonomously (approving loans, flagging transactions, processing claims) and only escalate exceptions to humans." That shift changes everything about governance. When humans make decisions, you can audit the human. When AI makes decisions at scale, you need to audit the entire system: the data, the model, the context, the logic, the exceptions, in real time.

These three shifts converge into a single imperative: financial institutions must build AI systems that are not just intelligent, but institutionally trustworthy.

THE REAL BOTTLENECK: WHY PROOFS OF CONCEPT DIE IN RISK COMMITTEES

Here's the pattern Foley has observed across dozens of financial institutions: a talented data science team builds an AI model that performs brilliantly in testing. Fraud detection accuracy improves by 30%. Credit decisioning time drops from days to minutes. The business case is compelling, the technology works, and everyone's excited. Then the model goes to the risk committee.

"That's where things fall apart," Foley says. "Because suddenly you're not talking about model accuracy. You're talking about explainability. You're talking about data lineage. You're talking about what happens when a customer challenges a decision, or when a regulator asks you to demonstrate that your model isn't discriminating against protected classes, or when the model produces an outcome that contradicts fifty years of institutional credit expertise."

The problem isn't that the models are wrong. The problem is that institutions can't demonstrate, with the rigour that regulators and risk committees demand, that the models are right for the right reasons. They can't trace decisions back to source data. They can't explain why the model weighted certain factors over others. They can't demonstrate that the training data was representative, unbiased, and properly governed. They can't show that the model will behave predictably when market conditions change.

"In financial services, AI decisions have direct human consequences," Foley emphasises. "If your credit model denies someone a mortgage, that's not an abstract algorithmic output. That's a family that can't buy a home. If your fraud detection system flags a legitimate transaction, that's a customer stranded abroad without access to their money. If your claims processing agent rejects a claim, that's someone facing financial hardship. Regulators understand this. They're not being obstructionist when they demand explainability and auditability. They're doing their job."

The bottleneck, then, isn't building AI that works. It's building AI that can be defended: at scale, at speed, under scrutiny, three years after deployment when the environment has changed, and the original data scientists have moved on.

WHY QUARTERLY GOVERNANCE CANNOT GOVERN MILLISECOND DECISIONS

Traditional governance models were designed for a different era. Monthly risk committee meetings. Quarterly policy reviews. Annual audits. Humans making decisions that could be documented, reviewed, and challenged through established processes.



"That model breaks down completely when you're dealing with AI agents making thousands of decisions per second," Foley explains. "If your AI is approving credit applications in milliseconds, you can't have a human review every decision. If your fraud detection system is flagging transactions in real time, you can't wait for the next governance committee meeting to validate the logic."

The regulatory evidence is mounting. The Financial Conduct Authority has issued warnings about AI governance failures. Consent orders have been issued against institutions that couldn't demonstrate adequate oversight of automated decision-making. The EU AI Act explicitly requires real-time monitoring and auditability for high-risk AI systems, which includes most financial services applications.

"The paradigm has to shift," Foley argues. "Governance can't be retrospective anymore. It can't be something you do after the fact, reviewing decisions that have already been made. It has to be embedded, applied at the point of data consumption, at the moment of decision-making, with context captured in real time so that you can reconstruct, investigate, and demonstrate compliance when required."

This is where most organisations struggle. Their governance frameworks are built around policy documents, approval workflows, and human oversight. They're not built around real-time enforcement, automated auditability, and semantic understanding of what data means in business context.

"If you're still thinking about governance as a quarterly review process, you've already lost," Foley says bluntly. "You need governance as architecture:

embedded in the systems, enforced at the point of consumption, capturing context automatically so that every decision is auditable by design."

THE BUSINESS MEANING LAYER: WHY METADATA MATTERS MORE THAN MODELS

Here's a problem that doesn't get enough attention: AI models are extraordinarily good at finding patterns in data. They're extraordinarily bad at understanding what that data actually means.

"Models understand correlation, not causation," Foley explains. "They understand statistical relationships, not business logic. And when you feed them data that lacks semantic context (when the same metric is calculated differently across different systems, when 'customer' means something different in your CRM than it does in your core banking platform, when you've got inconsistent definitions left over from three different acquisitions), the model will confidently produce answers that are technically correct but semantically wrong."

This is the root cause of many AI hallucinations in enterprise settings. The model isn't malfunctioning. It's doing exactly what it's trained to do: finding patterns in noisy, inconsistent, semantically incoherent data and making inferences that seem plausible but are actually nonsense.

The solution, Foley argues, is metadata, a business meaning layer that sits above the physical data and provides context, definitions, business rules, and governance policies. "Metadata tells you what 'customer' means across your entire organisation," he explains. "It tells you how 'revenue' should be calculated, who owns this data, where it came from, what transformations have been applied, what

governance policies apply, and what business rules constrain how it can be used."

This isn't just about data quality. It's about making governance scalable. "If you're writing governance policies at the table level (this specific table in this specific database has these specific rules), you'll never scale," Foley says. "But if you're writing policies at the semantic level (all data about customers, regardless of where it's stored, must comply with GDPR, must have consent recorded, must have lineage tracked), suddenly governance becomes architecture, not documentation." We successfully implemented this in one of our Financial Services clients where we enabled standardization across multiple geographies enforcing policy and meeting each country's requirements which in turn for example, enabled them to move AI models between countries in days vs months"

Foley highlights a direct operational consequence: "AI agents that consume data without governed semantic context are far more likely to hallucinate — to produce outputs that are statistically confident but factually or contextually wrong — because they are reasoning about data whose meaning they have had to infer rather than receive. Governed metadata eliminates that inference. It gives the AI the same grounded business context a senior analyst brings to a decision. That is not a nice-to-have — it is what separates AI that can be trusted and scaled from AI that remains permanently in the pilot phase.

According to data from Gartner, 60% of AI projects will be abandoned through 2026 due to poor data readiness — and for those that progress beyond POC before failing, the cost of each stalled deployment runs well into the hundreds

of thousands before opportunity cost is even considered. The question senior executives and boards should be asking is not what AI costs to build, but what ungoverned AI initiatives cost to abandon."

And this, in Foley's view, is where the conversation about metadata and semantics becomes not a technical discussion but a strategic one.

WHAT DO YOU ACTUALLY DIFFERENTIATE ON?

The build-versus-buy debate in enterprise technology is as old as enterprise technology itself. But Foley thinks the question is framed wrong.

"The better question is: where do you differentiate versus where do you commoditise?" he suggests. "Because if you're a bank or an insurer, your competitive advantage isn't in building semantic layers or metadata catalogues or data virtualisation platforms. Your competitive advantage is in fifty years of accumulated credit expertise. It's in your actuarial models. It's in your underwriting methodologies. It's in your customer relationships and your understanding of risk in specific markets and your ability to interpret regulatory requirements in ways that balance compliance with commercial opportunity."

That institutional knowledge, that domain expertise, is irreplaceable. Infrastructure, by contrast, is increasingly commodity. "Every bank needs semantic governance. Every insurer needs metadata management. Every financial institution needs data lineage and auditability," Foley observes.

These capabilities are foundational and universally necessary; building and maintaining it yourself is



not where a financial institution creates competitive differentiation, you're essentially paying a strategic tax: absorbing engineering talent that should be focused on the things that actually differentiate you - it takes two to three years at best when you could be operating in weeks, and it creates lock-in to your own technical debt rather than freeing you from anyone else's.

Reconciling and preparing data is the hidden tax on every AI programme we encounter in financial services. We have seen this working with one of our financial services clients this year, where AI-driven automation is now automatically reconciling and remediating over 90% of data quality issues that previously consumed significant engineering and business resource. Freeing that capacity for higher-value work has accelerated project timelines by more than seven months and liberating significant capital release within the same financial year — that is what governed, automated data preparation delivers to AI driven projects when it replaces manual reconciliation as the default."

This is where specialist platforms create value. "If you can adopt a platform that gives you semantic governance, metadata management, and data virtualisation without requiring you to rip out your existing stack, you've just freed up your engineering team to focus on building institutionally intelligent AI: the stuff that makes your accumulated knowledge computable, that enables your AI agents to make decisions that reflect your risk culture and your customer relationships and your regulatory interpretations - ie your decision intelligence which is your competitive edge."

Foley is careful to distinguish this from outsourcing strategic capability. "You're not outsourcing decision-intelligence. You're not outsourcing your models or your business logic or your customer

intelligence. You're adopting infrastructure that lets you focus on those things. It's the difference between building your own data centre and using cloud infrastructure. The cloud doesn't make you less differentiated. It makes you more focused on the things that actually create competitive advantage." In banking: the credit models tuned to your customer base, the risk frameworks calibrated to your portfolio, the relationship intelligence accumulated over years of client history. In insurance: the actuarial models refined across decades of loss data, the underwriting rules that reflect your appetite and your market, the claims intelligence that separates accurate reserving from expensive guesswork. And across both: the AI agents trained on your own regulatory interpretations, your own risk culture, your own definition of what a good decision looks like. That is where your edge is — and that is precisely what a specialist platform such as Stratio Decision Intelligence should free you up to focus on, rather than distract you from.

Stratio Decision Intelligence leverages your existing enterprise systems. It connects avoiding rip and replace enabling the governed business meaning layer to provide the necessary business context for AI/Agents to consume trusted information. Bringing all of this together means decisions and actions can be automated with AI, scaled and guarantee the audit and traceability Financial Services organizations need. Weeks vs months or years.

Speed matters enormously right now. The window for competitive advantage in AI is closing faster than most boards appreciate. Organisations that spend two years building what they could procure in weeks are not being strategic — they're being slow."

DATA SOVEREIGNTY AND THE LOCK-IN PROBLEM

There's a pattern that plays out repeatedly in enterprise technology: an organisation commits to a single hyperscaler, or a single platform vendor, or a single technology stack. The initial benefits are real: faster deployment, integrated tooling, vendor support. But five years later, the organisation discovers it's locked in. It can't easily migrate to alternatives. It can't audit what's happening inside the black box. It can't operate independently if the vendor relationship sours or the commercial terms become unfavourable.

"We've seen this movie before," Foley says. "And now it's not just a commercial problem. It's a regulatory problem."

He's referring to the wave of European regulation focused on operational resilience and data sovereignty. DORA, the Digital Operational Resilience Act, explicitly requires financial institutions to ensure they can operate independently of critical third-party providers. The EU AI Act requires organisations to maintain oversight and auditability of AI systems, even when those systems are provided by external vendors. Data residency requirements are tightening. Regulators are increasingly sceptical of arrangements where institutions can't fully account for how their data is used or how their AI systems make decisions.

"Interoperability isn't just a nice-to-have anymore," Foley argues. "It's sovereignty. It's the ability to own your data, own your decisions, and demonstrate to regulators that you're in control of your technology stack rather than being controlled by it."

This is where Foley sees Stratio's approach as differentiated vs other solutions that require centralization or replication of data. "We're not asking organisations to rip out their data lakes or replace their warehouses or abandon their BI tools, further increasing the data architecture inefficiency through duplications and writing off previous investments which are already increasing existing cost bases by at least 20%" he explains. "We're providing a semantic control plane that sits across your existing investments and gives you governance, metadata management, and decision intelligence without creating new dependencies or requiring data replication. You maintain sovereignty. You maintain interoperability. You reduce your cost base. And you can demonstrate to regulators that you're in control."

It's a philosophy born from hard experience. "The organisations that will thrive in the next decade are the ones building on open, interoperable foundations today," Foley predicts. "The ones that are locked into proprietary stacks are going to find themselves increasingly constrained, not just commercially, but from a regulatory perspective as well." Foley also highlights that whilst our conversation is focused on Financial Services, this is equally true across any regulated industry from Energy and Pharma through to Retail and the Public Sector.

READY TO SCALE: BANKING AND INSURANCE

So where are the real opportunities? Where can financial institutions deploy AI with confidence today, and where do they still need to build foundations?



Behind the Solutions

In banking, Foley sees a clear distinction between ready-to-scale and frontier opportunities. "Ready to scale right now are the high-frequency, rule-bound, data-rich decisions where governance requirements are well understood, and ROI is clear," he says. "Fraud detection and prevention. Credit scoring with explainable outputs. AML and KYC automation with full audit trails. Regulatory reporting with complete lineage. These are use cases where the data exists, the rules are established, and the governance frameworks are mature enough to satisfy regulators."

The frontier (three to five years out) is what Foley calls "institutionally intelligent AI." "Imagine AI that makes your bank's accumulated knowledge computable," he suggests. "Fifty years of credit expertise. Decades of regulatory interpretations. Deep relationship intelligence about customers and their businesses. That's the AI that enables relationship managers to have every piece of context about a customer instantly available, that lets compliance officers query decades of precedent in natural language, that allows credit committees to make decisions informed by institutional memory rather than just recent data."

The constraint isn't technology. "The models exist. The compute exists. What doesn't exist yet, in most institutions, is the data and governance foundation that makes that AI trustworthy," Foley explains. "You can't make institutional knowledge computable if that knowledge is trapped in inconsistent systems, undocumented processes, and the heads of people who are retiring."

In insurance, the opportunities come at different angles. "Claims processing is extraordinarily compelling," Foley says. "An AI agent that can

ingest a claim, validate it against policy terms, cross-reference loss history, assess fraud risk, and produce a recommendation with full audit trail: that compresses weeks into hours, sometimes minutes. And because insurance claims are highly structured and rule-bound, the governance requirements are manageable."

Underwriting is another frontier. "Dynamic risk assessment that draws on structured and unstructured data (customer surveys, satellite imagery, IoT sensor data from vehicles or properties) and produces explainable pricing decisions. Solvency II and the EU AI Act are actually accelerating this because they're forcing insurers to demonstrate that their risk models are explainable and auditable."

The three-to-five-year horizon in insurance, Foley believes, is genuinely real-time dynamic risk modelling. "Not models trained on historical data and updated quarterly. Models that reflect current reality (current weather patterns, current economic conditions, current customer behaviour) and adjust pricing and risk assessment in real time. Real-time risk modelling means you know who the individual customer actually is today — not who they were when they signed the policy — and that is what makes hyperpersonalised pricing and coverage possible. That's transformative. But it requires data infrastructure and governance that most insurers don't have yet."

There's a lesson that cuts across both sectors. "Insurance has a discipline around uncertainty and scenario modelling that banking could learn from," Foley observes. "Actuarial rigour is fundamentally about quantifying uncertainty and making decisions under ambiguity. That's exactly what AI governance

requires. Banks tend to think in terms of rules and thresholds. Insurers think in terms of probability distributions and confidence intervals. That mindset is incredibly valuable when you're trying to govern AI systems that are inherently probabilistic."

THE FINAL REFRAMING: GOVERNANCE AS PRECONDITION, NOT CONSTRAINT

We're sitting in Stratio's offices as our conversation winds down, and I ask Foley what he'd want CIOs, CDOs, and AI leaders to take away from this conversation. His answer is immediate.

"Stop thinking about AI governance as a constraint on innovation," he says. "Reframe it as a precondition for scalable deployment. Because the critical question isn't 'Can AI do this?' The critical question is 'Can we defend this decision to a regulator, a customer, or the board in three years when this AI has been running at scale, and the environment has changed?'"

If the answer to that second question is no, Foley argues, the first question becomes irrelevant. "Trusted AI isn't slower AI. It's AI that enables you to scale because it's got the foundations to survive scrutiny, to be audited, to earn the confidence of the stakeholders who need to say yes. Organisations building weak foundations aren't moving faster. They're just accumulating technical debt, regulatory debt, and confidence debt."

It's a powerful reframing. In an industry obsessed with speed and innovation, Foley is arguing that the fastest path forward runs through governance, not around it. "Governance isn't a layer you add to AI after the fact," he insists. "It's the architecture

on which AI operates. It's embedded in the data, enforced at the point of consumption, captured in metadata, and validated in real time. That's what compliant AI for enterprise decisions really means. And it's the only AI that's going to scale."

As I leave, I'm struck by how counter-intuitive Foley's thesis is, and how obviously correct it becomes once you see it. The institutions racing ahead with ungoverned AI aren't actually ahead. They're building systems that will fail the first time they face serious scrutiny. The institutions that seem to be moving slowly, building foundations and embedding governance, are actually building the only AI that will survive contact with regulators, customers, and reality.

In financial services, trust isn't a nice-to-have. It's the product. And AI that can't be trusted can't be sold, can't be scaled, and ultimately can't survive. The organisations that understand this, that build governance as architecture rather than afterthought, won't just deploy AI faster. They'll redefine what competitive advantage means in an industry where decisions are increasingly made by machines that must be, above all else, defensible.

That's not a constraint on innovation. That's the foundation on which real innovation becomes possible.





Compliant AI for Enterprise Decisions

AI can only act with precision when meaning, context and control come first. Speed without governance isn't an advantage, it's a decision you can't defend later.

Explore AI built for enterprise control

stratio.ai