

RE • WORK

15 SEPTEMBER 2026

AI in Finance Summit London

The must-attend summit for solving the engineering and governance hurdles of actual AI deployment.

AI adoption in financial services is no longer the challenge; scaling it safely, reliably, and with measurable impact is. As AI in Finance returns to London for its 11th edition, it brings together the leaders and practitioners navigating this exact shift: **moving beyond pilots to production-ready systems that are tested, governed, and resilient by design.**

WHAT'S NEW IN 2026?

- **Execution Over Strategy.** A working-session format built around deployment problems, not AI theory.
- **The "Failure Files".** Candid post-mortems from real production failures at financial institutions.
- **End-to-End AI Security.** From non-human identities to AI supply chain attacks, the 2026 threat landscape.
- **Foundation Model Strategy.** Live debates on building, buying, or renting, and the economics behind each path.
- **Hands-On Workshops.** Stress-testing and adversarial AI sessions you won't find anywhere else.
- **EU AI Act Six Weeks In.** First post-deadline assessment and what's actually changed.
- **Cross-Event Networking.** Practitioner-led networking with experts across AI in finance and CDAO FSI events.

Key themes:

- From Pilot to Production
- Foundation Model Strategy
- AI Security and the Agent Attack Surface
- EU AI Act High-Risk Compliance
- Agentic AI Design and Control
- Model Risk in the Age of Agents
- AI in Capital Markets and Quant
- Red-Teaming and AI Validation
- AI Sovereignty and Vendor Concentration
- The CAIO Role in 2026–27

 **15th September 2026**  **America Square Conference Centre, 1 America Square, 17 Crosswall, London EC3N 2LB**

Questions about the programme? Get in touch with Conference Producer **Emilia Malaj** at emilia.malaj@coriniumgroup.com

Confirmed Speakers:

Roxanne Howdle-Rowe, *Managing Director, Group Head of Data & Analytics Office* — **BRITISH BUSINESS BANK**
Annapoorani Lakshmi Narayanan, *Product Manager Lead, AI Research* — **J.P. MORGAN CHASE**
Parag Kumar, *Executive Director, Data Analytics Lead, Compliance, Conduct and Operational Risk* — **J.P. MORGAN CHASE**
Jai Ferguson, *AI Regional Lead for Europe* — **HSBC**
Kaushik Chaubal, *Senior Director, Investor Solutions* — **BLACKROCK**
Hemant Patkar, *Cybersecurity Lead Designer/Architect* — **VIRGIN MONEY**
Parinita Kothari, *Engineering Lead* — **LLOYDS BANKING GROUP**
Charles Phiri, *Executive Director, SME AI/ML Innovation* — **J.P. MORGAN CHASE**
Eshaan Salwan, *IT Auditor* — **GOLDMAN SACHS**
Soung Low, *Model Risk Data Scientist* — **NATWEST GROUP**
Amit Kumar, *Associate Director, eFx Quant* — **HSBC**
Rajasree R, *Product Lead - Agentic Observability* — **LLOYDS BANKING GROUP**
Kirsty Montignani, *Head of Security Data & AI* — **LLOYDS BANKING GROUP**
Sebastian Obeta, *Member of European AI Alliance* — **EUROPEAN COMMISSION**
Temi Afeye, *Technical Lead, Senior AI Scientist* — **LLOYDS BANKING GROUP**
Dr David, Crelley, *Head of Responsible AI and Data* — **ADMIRAL GROUP PLC**

Pending Speakers:

Maninder Kartik, *CIO - Customer & Marketing* — **AVIVA**
Miles Hillier, *Chief Product & Technology Officer, AI* — **NATWEST GROUP**
Sabina Kuzina, *Data Science Platform Lead & AI Training Lead* — **DEUTSCHE BANK**
Arunita Roy, *Technical and Data Strategy Lead, Transforming Data Collections Programme* — **FCA**
Pilar Roig, *VP, Real-Time Payments Data Product* — **MASTERCARD**
Val Riabtsev, *Head of Platforms, AI Agents & Automation* — **REVOLUT**
Alok Kulkarni, *Director, Conversational Banking, Engineering and AI* — **VIRGIN MONEY**
Harshit Singhvi, *Senior Vice President, Generative AI* — **CITI**
Emily Campbell-Ratcliffe, *Head of AI Assurance, Responsible Technology Adoption Unit* — **DEPARTMENT FOR SCIENCE, INNOVATION AND TECHNOLOGY**

AI in Finance London 2026 | Programme

15 September 2026

8:00-
8:45

Registration and Coffee in the Exhibition Area

8:50-
9:00

Chairperson's Opening Remarks and Icebreaker

9:00-
9:30

OPENING PANEL: The Board Has Approved the Budget. Now What? A nHonest Assessment of Where AI Is Actually Delivering

- Real-world scale outcomes vs. the business cases that funded them - which financial services use cases have genuinely reached industrial production in 2026?
- The "deployment wall" and what are the specific organisational, technical, and governance triggers that cause promising AI programs to stall between PoC and scale?
- Where is the gap structural (data, talent, legacy) vs. closeable (tooling, process, governance)? Digital-native vs. incumbent differentiation
- What does "industrial-grade AI" require, and why is it a different engineering problem than most firms anticipated?

Panellists:

Roxanne Howdle-Rowe, *Managing Director, Group Head of Data & Analytics Office* — **BRITISH BUSINESS BANK**
Parag Kumar, *Executive Director, Data Analytics Lead, Compliance, Conduct and Operational Risk* — **J.P. MORGAN CHASE**
Kirsty Montignani, *Head of Security Data & AI* — **LLOYDS BANKING GROUP**
Moderator: Jai Ferguson, *AI Regional Lead for Europe* — **HSBC**

9:30-10:00

KEYNOTE: Sovereign AI in the Enterprise: De-risk, Deploy and Scale Agents with VMware Tanzu

Every day brings promises of a new, transformational agentic use case. Meanwhile, your organization may be struggling to move from proof-of-concepts to secure, production-grade agents that you can govern. Achieving this requires a strategic, enterprise-wide approach to operationalizing and securing AI and the data it needs to be useful. In this session, Purnima Padmanabhan introduces an actionable framework designed to address the real-world challenges of securing and operationalizing agents in highly regulated industries like Finance. We will explore how to maintain compliance, sovereignty and security velocity covering topics such as:

- Agentic Runtime Considerations: Establishing guardrails, observability, and standardization for autonomous AI agents to ensure they operate within secure, compliant perimeters.
- Data & AI Sovereignty: Balancing regulatory compliance and local control with the flexibility of an open ecosystem of models, frameworks and data sources.
- Robust, Centralized Governance: Enforcing strict controls on agent access to sensitive enterprise tools, data, and infrastructure while maintaining sovereignty.
- Defending Against AI-Enabled Threats: Strategies for compressing the vulnerability-to-patch cycle to ensure your software supply chain is protected against the growing threat of AI-driven security vulnerabilities.

Join us to learn how to operationalise agents without compromising the integrity or security of your data and enterprise systems.

Purnima Padmanabhan, *General Manager, Tanzu Division*

BROADCOM

10:00-10:30

FIRESIDE CHAT: Six Weeks In — What the EU AI Act High-Risk Obligations Have Broken, Fixed, and Left Unresolved

- Which financial services AI systems are now classified as High-Risk under the EU AI Act (post-August 2026), and what specific documentation, testing, and human oversight requirements are now mandatory?
- What has been harder than expected? Where compliance programs have encountered unexpected technical, legal, and operational friction
- What has been easier than expected? Practical shortcuts and interpretations that have emerged from early NANDO notified body assessments
- Which UK firms operating under FCA/PRA supervision are now misaligned between EU AI Act obligations and domestic SS1/23 model risk requirements, and how to resolve it

Sebastian Obeta, *Member of European AI Alliance* — **EUROPEAN COMMISSION**

Dr David, Crelley, *Head of Responsible AI and Data* — **ADMIRAL GROUP PLC**

10:30 – 11:00

Mid-Morning Coffee Break & Networking in Exhibition Area

11:00-11:20

TECHNICAL KEYNOTE: Owing Your AI Stack: Infrastructure Architecture Decisions That Separate Production Systems from Expensive Prototypes

- The production AI system architectural blueprint - what components must be owned vs. procured, and what are the vendor lock-in risks in each layer?
- The hardest infrastructure problems - latency management in multi-step LLM chains, state management across agent sessions, and consistency in high-throughput inference
- How to implement continuous LLM evaluation against financial reasoning, regulatory compliance, and domain accuracy benchmarks?
- What does production observability look like when your AI system is making thousands of consequential decisions per hour?

11:20-11:40

Building Trustworthy AI in Financial Services: Security, Resilience and Governance Beyond Prompt Injection

- Defending against "excessive agency" and unauthorised actions as autonomous systems become primary targets for cybercrime.
- Moving beyond system prompts to secure the non-human identities and "digital insiders" who manage high-value transactions.
- Strategies for securing the AI supply chain—from data poisoning and model tampering to deepfake-driven social engineering.
- Building "Defensible AI" that maintains operational integrity under the scrutiny of the EU AI Act and DORA.

Hemant Patkar, *Cybersecurity Lead Designer/Architect*

	VIRGIN MONEY	
11:40-12:10	PANEL: From Shadow AI to Automated Guardrails: How Model Risk and Engineering Are Finally Solving Governance Together <ul style="list-style-type: none"> • The shadow AI crisis - how widespread is unsanctioned LLM usage in UK/EU financial institutions, and what are the realistic options for discovery, containment, and governance? • Moving from abstract responsible AI principles to engineering-enforced controls looking at a production ML pipeline. • EU AI Act High-Risk classification - which financial services use cases are affected, and what documentation and human oversight requirements are now mandatory? • When an agentic AI system causes a regulatory breach, who is responsible? The model owner, the data team, the business, or the vendor? <p>Panellists: Soung Low, <i>Model Risk Data Scientist</i> — NATWEST GROUP Annapoorani Lakshmi Narayanan, <i>Product Manager Lead, AI Research</i> — J.P. MORGAN CHASE Eshaan Salwan, <i>IT Auditor</i> — GOLDMAN SACHS Moderator: Temi Afeye, <i>Technical Lead, Senior AI Scientist</i> — LLOYDS BANKING GROUP</p>	
	USE CASE SHOWCASE 	
12:10-12:20	Showcase 1: Real-Time AI Decision Explainability in Fraud Prevention Under Consumer Duty <ul style="list-style-type: none"> • How AI fraud detection is deployed in a live UK payments environment, with real performance metrics • What happens when your AI model flags a legitimate customer and you need to explain the decision • How is explainability tooling being integrated into production fraud workflows, not added as an afterthought 	
12:20-12:30	Showcase 2: When Your AI Agent Gives Bad Advice — Managing Hallucination, Escalation, and Consumer Duty Liability in Customer-Facing Deployments <ul style="list-style-type: none"> • What happens when an AI customer agent provides financially consequential incorrect information, and how is it managed operationally and legally? • How are firms documenting AI-assisted customer interactions to demonstrate fair outcomes and appropriate escalation? • The emerging standard for human-in-the-loop escalation design in FCA-regulated customer-facing AI 	
12:30-12:40	Showcase 3: Beyond Basic RAG — Knowledge Graphs, Structured Retrieval, and Solving the Accuracy Problem in Enterprise AI <ul style="list-style-type: none"> • Why vanilla RAG fails in regulated financial contexts: hallucination rates, citation reliability, and the accuracy floor problem for compliance and risk applications • Knowledge graph-augmented retrieval in practice: architecture decisions, implementation complexity, and measurable accuracy improvements • How to evaluate and benchmark RAG performance in financial services — the metrics that matter and those that mislead 	
12:40-13:40	<i>Lunch and Networking in the Exhibition Area</i>	
	MAIN CONFERENCE Workshops	TRACK B Roundtables
13:40-14:20	WORKSHOP A: Engineering Accountability — Designing Controlled Workflows for Agentic AI <ul style="list-style-type: none"> • Identifying the "Goldilocks" use cases for agents: tasks with high complexity but clear, programmable constraints. • How do you design workflows with hard-coded "Circuit Breakers" and human-intervention triggers? • Frameworks to ensure consistency across agent systems and meet internal audit and reliability standards. 	<p><i>Two simultaneous roundtables running in parallel. Attendees remain seated at their chosen table throughout. Speakers rotate between tables at 40-minute mark, ensuring both groups hear from every facilitator.</i></p> <p><i>Each table will have a printed prompt card with 3 discussion questions to guide conversation and encourage audience participation.</i></p> <p><u>Human + AI Operating Model and Adoption at Scale</u></p>

	<ul style="list-style-type: none"> How to mitigate the risks of data leakage and "permission creep" as agents navigate internal banking silos? <p>Parinita Kothari, <i>Engineering Lead</i> — LLOYDS BANKING GROUP Rajasree R, <i>Product Lead - Agentic Observability</i> — LLOYDS BANKING GROUP</p>	<p>ROUNDTABLE 1: From Mandate to Habit — Why AI Adoption Is an Organisational Design Problem, Not a Technology One</p> <ul style="list-style-type: none"> Where are AI tools being used daily vs. abandoned after rollout—and why? How do you redesign workflows (not just tools) to ensure AI becomes part of decision-making? What incentives, training, or friction points determine whether teams adopt or reject AI?
<p>14:20-15:00</p>	<p>WORKSHOP B: The Pre-Flight Checklist for Stress-Testing and Validating Financial AI</p> <ul style="list-style-type: none"> Moving beyond "accuracy" to define rigorous KPIs for hallucination, bias, and financial reasoning. Practical frameworks for uncovering edge cases and security vulnerabilities before they reach production. Identifying intervention points where human oversight must be implemented into the automated workflow. How to monitor and improve performance to prevent model drift post-deployment? <p>Senior Representative, IBM</p>	<p>ROUNDTABLE 2: The £100m Mistake: What C-Suite Leaders Learned When Their AI Got It Catastrophically Wrong</p> <ul style="list-style-type: none"> What were the early warning signs that were missed? How are governance frameworks evolving to assign personal accountability for AI model failures to named executives? Who bears financial liability when an AI system causes material customer harm or regulatory breach?
<p>15:00-15:30</p>	<p><i>Afternoon Coffee Break and Networking in Exhibition Area</i></p>	
<p>15:30-16:00</p>	<p>FIRESIDE CHAT: Algorithmic Alpha vs. Systemic Stability — How Portfolio Leaders Are Navigating AI-Driven Alpha, Herding Risk, and Regulatory Scrutiny</p> <ul style="list-style-type: none"> How is AI fundamentally altering alpha generation, execution logic, and portfolio construction in 2026? Exploring the risks of "Model Herding"—could synchronised AI decision-making trigger new forms of market volatility and flash-instability? Where should firms draw the line between machine execution and human judgment? How are regulators approaching AI in capital markets, and what should firms prepare for next? <p>Panellists: Amit Kumar, <i>Associate Director, eFx Quant</i> — HSBC Moderator: Kaushik Chaubal, <i>Senior Director, Investor Solutions</i> — BLACKROCK</p>	
<p>16:00-16:30</p>	<p>KEYNOTE: The CAIO in the Crossfire — Authority, Budget, and the Limits of the Chief AI Officer Role in 2026</p> <ul style="list-style-type: none"> A candid assessment of where the role has real authority and where it is advisory or ceremonial Critical decisions for CAIO in 2026 -> build-vs-buy model strategy, EU AI Act compliance program ownership, and agentic AI governance framework How does the CAIO role intersect with (and conflict with) the CDO, CTO, CRO, and CPO, and how are high-functioning firms resolving these tensions? How will the CAIO title develop in the future? The skills gap, the regulatory exposure, and the board relationship that will define the role's next evolution 	
<p>16:30-17:00</p>	<p>CLOSING PANEL: Who Is Actually in Control of Your Firm's Intelligence in 2026? Concentration, Sovereignty, and the Accountability Mandate</p> <ul style="list-style-type: none"> When three hyperscalers and two foundation model vendors control the inference layer of the global financial system, what is the realistic failure scenario, and what are regulators doing about it? What does a "demonstrable operational control" mean when your core AI capability is a third-party API? A look at specific architectural and contractual requirements that satisfy regulators What compliance now requires for production AI systems, which firms are unprepared, and what the enforcement timeline looks like How leading firms are managing the strategic tension between best-in-class third-party AI and the obligation to maintain explainability, control, and exit optionality 	

17:00-
17:10

Chairperson's Closing Remarks

17:10-
18:00

Networking Drinks Reception and Prize Draw

END OF SUMMIT