

The Architecture of Trust in AI-Driven Banking

Banking has been one of the biggest spenders when it comes to investments in AI. Yet research shows that while banks are accelerating adoption, most are deploying without adequate oversight or governance, a dangerous condition for a sector that depends on trust more than almost any other. As per the SAS IDC study Data and AI Impact Report: The Trust Imperative, only 11% of banks have achieved internal confidence in AI alongside systems that are demonstrably trustworthy. The same study found that nearly half of banks are caught in what IDC calls the “trust dilemma”, either underusing AI they have validated or over relying on systems that have not been adequately tested. For a sector under continuous scrutiny from regulators, customers, and investors, this is not a tolerable condition.

“Trust in AI-first banking is not a focus area or a commitment statement. It needs to be engineered, layer by layer.”



Achieving this requires a partner with both the orientation and the depth to work across four interconnected areas.



AI at the Core, Not the Edges

Embedding AI into core systems, operations, and decision layers - moving banks from isolated pilots to enterprise-scale transformation.



Principles that Engineer Trust

Embedding fairness, explainability, privacy, reliability, and regulatory compliance into every AI-led system by design.




Pragmatic, Outcome-Driven Approach

Delivering measurable business outcomes across onboarding, decisioning, fraud, modernization, and operational efficiency.



Deep Banking DNA

25+ years of exclusive banking focus, ensuring AI is aligned to real-world banking operations, regulations, and customer realities.



AI at the Core, Not at the Edges

A central reason for the trust gap is that most AI in banking is still deployed at the periphery or as isolated use cases. Evident AI Index data indicates that half the AI tool announcements made by banks in the first quarter of 2026 were not in production. The same research points out that the share of new use cases in the testing or pilot stage announced by banks tripled compared to the previous year.

When AI operates in pilots, risks are contained and results are more predictable. When applied at the core, the landscape changes materially, production data variability introduces noise that was absent in controlled environments, integration with live systems surfaces dependencies that were invisible at the pilot stage, and governance that seemed manageable during testing becomes a binding constraint at scale.

Moving AI to the core is not a matter of flipping a switch. It involves four fundamental architectural shifts: systems of record becoming systems of intelligence; batch processing giving way to event-driven architecture; monolithic systems evolving into composable architecture; and static rules being replaced by adaptive decisioning. Each shift carries its own governance mandates and risk windows that must be designed for, not retrofitted.

Guaranteeing trust in AI-first banking at this level requires a four-layer trust architecture:



At the pilot stage, gaps in any of these layers are largely inconsequential, the blast radius is limited. When AI is deployed at the core, each layer becomes load bearing. A data trust failure, for example, means an AI fraud detection model trained on historical patterns systematically misclassifies behavior that has shifted since training. A model trust failure means a credit decision triggers a regulatory inquiry because no audit trail supports the inference. These are not hypothetical risks; they are the governance failures that have already prompted regulatory action in multiple jurisdictions.

Maveric embeds AI into core systems, operations, and decision layers, not at the periphery through isolated pilots. This shifts banks from fragmented experimentation to enterprise-wide AI-first models with real governance, observability, and accountability.

“Architectural depth at the core must, however, be matched by design principles that make scale safe. That is the work of the third strength.”

Principles That Engineer Trust


Most AI governance in banking today is retrospective - audited after deployment, documented after an incident, explained after a regulatory query. This is not governance; it is remediation. Compliant by design, not by audit is the architectural alternative: embedding governance into the system from the beginning rather than appending it after the fact.

In practice, this means that the four principles of fairness, explainability, reliability, and privacy are specified before model training begins, not assessed afterwards. Audit trails are generated as a natural output of the delivery pipeline rather than reconstructed at the onset of a compliance event. The difference is structurally significant - reconstruction is approximation, whereas embedded traceability is evidence.

Trust principles should also aim at model accountability rather than model accuracy. Accuracy is hygiene, it describes how often a model is right in a controlled environment. Model accountability means being able to explain, justify, and defend each decision the model takes. This requires designing AI so that every decision can be traced back to data and logic, explained in terms that are meaningful to regulators and customers alike, and continuously monitored for drift and bias.

The shift from accuracy to accountability is not semantic; it changes what gets built, how it gets tested, and what constitutes a deployment-ready system.

Maveric applies a standards-led engineering approach - embedding the principles of fairness, explainability, reliability, and privacy alongside global banking standards and regulatory requirements from the architecture stage. This includes



“Principles without measurable outcomes, though, remain philosophy. The fourth strength is what closes the loop between safe AI and valuable AI.”

Pragmatic and Outcome Driven

Most AI initiatives operating at a pilot level are not connected to measurable business outcomes and that is precisely why they struggle to earn sustained business buy-in. As per Evident, only 15 of the 50 banks it tracks share projected or realized gains from AI. ROI is not validated at the production stage, attribution frameworks are absent, and output anchors are not established to measure productivity gains. Without these, there is no mechanism to connect AI activity to business impact.

A significant structural reason for this gap is that AI use cases are typically scoped by technology teams without a corresponding business owner who is accountable for the outcome. This severs the accountability chain that makes impact measurable. The predictable result is an organization that tracks activity metrics - adoption milestones, use cases in pilot and production, cost benefits observed in controlled tests, while the value gap widens, because these measure deployment, not impact.

Moving from isolated use cases to organization-wide adoption starts with business revisioning:



These three elements — output anchors, attribution frameworks, and outcome chains — must be established before development begins, not after. They are the architecture of accountability that makes AI adoption irreversible rather than episodic.

The SAS IDC study found that AI investments focused on customer experience generated a return of \$1.83 for every dollar invested, above cost-reduction-led initiatives at \$1.54. The implication is direct: outcome-first scoping does not just create accountability; it creates better returns.

At Maveric, every engagement is anchored to quantifiable business impact - faster customer onboarding, real-time decisioning, reduced cost-to-serve, improved fraud detection, and accelerated modernization.

“Outcome-driven execution is how trust engineering translates into the business results that banking leaders are accountable for.”

Deep Banking DNA

Domain depth is not a superficial or theoretical knowledge of banking products. It is an in-depth understanding of banking products, processes, workflows, integrations, and regulations in real-world conditions and this is absolutely critical when developing AI applications. Consider a credit workflow application. Without understanding the lineage of data that feeds the system, the implications on an inference engine that connects to a multitude of upstream applications cannot be adequately assessed. Errors in data lineage do not surface during development; they surface in production, under regulatory examination, or in a customer-facing decision that cannot be explained.

Generic AI governance frameworks are unable to address the specificity demanded by banking regulations, whether BCBS 239, which requires demonstrable data lineage and accuracy under stress conditions; SR 11-7 (now updated to SR 26-2), which governs model risk management and mandates validation of every model before deployment, including AI models; or the EU AI Act, which classifies credit scoring and risk assessment tools as high-risk AI systems subject to pre-market conformity assessments. Under these frameworks, an AI credit scoring model is not just software - it is a regulated model. Engineering Trust means building AI that is compliant by design, which effectively mandates that those developing AI applications understand thoroughly what compliance means in practice for each use case and geography.

**Enterprise-wide or line-of-business-wide AI adoption
requires a framework that brings together**



**Safe-by-Design
Principles**



**Multi-Agent
Orchestration**



**Domain Data
Foundations**

A deep understanding of the banking context is non-negotiable for each element of this framework, without it, coherence across real banking workflows and regulatory obligations cannot be achieved. With over 25 years of exclusive focus on banking and financial services across Retail banking, corporate banking, wealth management, and capital markets, Maveric builds AI within the fabric of banking, not on top of it. Every AI model, use case, and solution is aligned to real-world banking processes, regulatory constraints, and customer realities across geographies.

**“Domain depth alone, however, only creates value when AI is
deployed where it matters most, at the operational core.”**

The Architecture of Trust

The four areas outlined above are not independent. Banking DNA provides the foundation - the context that is critical for AI to operate in alignment with real-world processes and regulatory obligations. AI at the core is the deployment architecture that makes trust engineering consequential, because trust frameworks that only govern pilots are not trust frameworks at all. Principles that engineer trust are the design constraints that make AI deployment safe to scale, turning governance from a retrofit into a structural property. Finally, a pragmatic and outcome-driven orientation creates the business case that sustains AI adoption beyond the pilot phase.

“The banks that win the next phase of AI-first banking will not be defined by who deployed AI first, whose models are most sophisticated, or who has the most comprehensive strategy. They will be defined by who built the infrastructure to trust their AI, in production, under examination, at enterprise scale, and over time.”

Authored by



P Venkatesh,
Co-founder and whole-time Director

As the Co-founder and whole-time Director at Maveric, P Venkatesh (PV) leads the global thought leadership function aimed at shaping and promoting Maveric's perspectives as well as expertise in the banking technology space. By building relationships with industry influencers, partners and BankTech ecosystem leaders, PV drives creation of impactful frameworks, methodologies and landscape reports that provide informed perspectives on new age technologies that shape the BankTech space.



Anoop Karumathil Melethil,
Vice President, Marketing

Anoop has two decades of experience in conceptualizing, creating, and driving technology-focused marketing initiatives across tier 1 as well as mid-size global IT services organizations. He has built his expertise in building marketing functions from the ground up, encompassing areas including but not limited to media, public relations, influencer marketing, digital marketing, employer branding, thought leadership, brand outreach, field marketing, ABM, and internal communication.

Maveric Systems

is a banking-exclusive technology specialist with over 25 years of domain expertise. We partner with global financial institutions to engineer trust in AI-first banking.

While others apply AI at the periphery, we embed it at the core through our proprietary AI @ Scale framework and AI-powered platforms and solutions. Guided by principles that engineer trust, we embed fairness, explainability, reliability, and compliance into every AI solution by design. This enables responsible AI adoption at scale.

Our domain depth across operations and technology in retail and corporate banking, wealth management, and capital markets, combined with a pragmatic, outcome-driven delivery model, ensures that every AI initiative is rooted in contextual relevance and precision.

Backed by dedicated AI Centers of Excellence, a powerful ecosystem of technology platforms, and recognition from leading industry bodies, we are the trusted engineering partner for the AI era.

The Maveric Edge

25+

years of domain mastery

7+

Operations and Technology AI CoEs

12+

AI powered Platforms and Solutions

AI@Scale

Scale Proprietary Framework for AI adoption at Scale

Maveric NXT Inc

5, Independence Way, Suite 300, Princeton, NJ 08540 USA

Reach us: marketing@maveric-systems.com

