

credx

Beyond the Credit Scorecard

How AI Is Reshaping Credit in the GCC



Foreword



Across Saudi Arabia and the wider GCC, the credit conversation is shifting from how much capital can be deployed to how intelligently it is allocated. The new infrastructure is data, models, and the institutional discipline to operate them responsibly.

Saudi AI commitments meet GCC fintech reality

\$14.9 B  **35%**
LEAP 2025 of GCC fintech

This article looks at where machine intelligence is delivering for credit in the GCC, what regulators expect of the institutions deploying it, and where human judgment still has to live in the loop.

Saudi Arabia's financial system is being rewired around data and computation.

In 2024, the Kingdom's locally managed financial assets crossed SAR 1 trillion (\$267 billion). Licensed fintech firms reached 261 against a 2030 target of 525, and 79% of retail transactions moved through digital rails, surpassing the 2025 target a year ahead of schedule.

At LEAP 2025 in Riyadh, \$14.9 billion in fresh AI investment was announced in a single event, with \$10.9 billion of that earmarked for infrastructure, AI initiatives, and startup funding. The artificial intelligence category already accounts for roughly 35% of GCC fintech activity, the largest of any technology category in the regional fintech stack.

What this means for credit is structural. Models drawing on alternative data, retrieval-augmented decisioning, and machine learning are extending credit to borrowers conventional underwriting cannot read. Agentic systems are taking over routine work in onboarding, screening, and monitoring.

The Credit Gap AI is Being Asked to Close

Saudi Arabia and its Gulf neighbors have an ambitious credit problem. Vision 2030 targets a 35% SME contribution to GDP by 2030, up from 21.9% in 2023. To reach that figure, financing has to flow to a long tail of small and mid-sized firms that the conventional banking system has historically struggled to serve.

The numbers tell the story. Globally, banks allocate roughly 22% of their lending to SMEs. In MENA, the figure falls below 10%. In the GCC, it sits under 2%, according to Deloitte. The result is a regional SME financing gap estimated at around \$250 billion. Banks frequently demand collateral coverage of 200 to 250%, a threshold that is unworkable for the asset-light, service-oriented businesses that now define a meaningful share of the non-oil economy.

Traditional credit assessment compounds the problem. It is built around historical financial statements, audited records, and tangible collateral. Many SMEs in the region present none of these in the form a conventional credit committee expects. Loan applications can take weeks; rejection rates in some segments are reported above 70%. The central question for any non-bank lender stepping into the gap is the same one that defeats the banks: how do you price risk on a borrower whose formal data does not describe the business well?



The next phase of GCC credit growth will depend on how institutions use data, technology, and risk infrastructure together.

How AI is expanding Credit Infrastructure

Models that can read alternative data, learn from observed behavior, and update continuously are moving the underwriting frontier inward, toward segments that were previously uneconomic to serve. Open banking has given them lawful, consent-based access to the data they need to do it.

SAMA's Open Banking Framework, launched as part of the Financial Sector Development Program, allows licensed third parties to read a customer's banking data with explicit consent. The framework now supports account information services and is gradually expanding into payment initiation, identity verification, and credit scoring. Combined with the National Strategy for Data and AI, it gives lenders both the raw inputs and the regulatory clarity to build modern underwriting on top.



Where formal credit histories are thin, lenders are turning to data the borrower already generates. Transaction streams from open banking APIs, supplier and buyer payment patterns from supply chain platforms, e-commerce sales velocity, telecom usage, utility records, customs declarations, point-of-sale activity, and e-invoicing data are each a credit signal the borrower's business has been broadcasting all along. The task is to read them well.

The combined effect is consequential. Roland Berger reports that automated credit models in the GCC now deliver real-time loan decisions, AI fraud detection has cut false-positive rates by up to 90%, and conversational systems resolve as much as 80% of customer queries without human intervention. Customer acquisition costs can fall by three-quarters, and compliance workloads by two-thirds, when the workflow is well designed.



Alternative data enhances credit assessment by providing broader operational visibility.

Machine learning models capture this complexity in ways linear scorecards cannot.

Gradient-boosting ensembles and modern neural architectures handle nonlinear interactions, tolerate the messy reality of alternative data, and produce continuous default probabilities that update with each new transaction. Anomaly-detection methods like autoencoders and isolation forests sit alongside them, flagging behavioral shifts that may signal financial distress weeks or months before a missed payment shows up in the books. Research on emerging market lending consistently finds that machine learning models outperform logistic regression on default prediction, particularly where labeled default data is sparse, which is exactly the GCC SME case.

Three Architectures Driving the Shift

Three distinct technical patterns are doing most of the work in modern credit, and each is different enough to be worth distinguishing. The first is machine learning, which sits closest to traditional risk modeling. Supervised models predict default probability from labeled historical outcomes; unsupervised methods detect anomalies and concept drift. In credit, this is where proprietary algorithms, trained on a lender's own portfolio data, deliver the most durable advantage. The signals they learn from are observable transactional behaviors: payment timing, balance volatility, deposit patterns, sectoral concentration. The second is **retrieval-augmented generation**, commonly abbreviated RAG, which is where large language models become useful for credit work without hallucinating. A RAG system pairs an LLM with a curated knowledge base, the lender's credit policy, sector reports, audited financial statements, regulatory filings, and grounds every output in retrieved source material. For underwriting, this means an analyst can request a structured assessment of a borrower's financials and industry context, and receive an answer that cites the exact documents it relied on. NVIDIA's 2024 financial services survey found 57% of firms already running RAG workloads in production.

The third pattern is agentic systems, the newest of the three and arguably the most operationally significant. Rather than answering a question, an agent executes a multi-step workflow. It pulls the documents, runs the screens, fills the forms, escalates the exceptions. McKinsey estimates 15 to 20% cost reduction across banking functions from agentic deployment. AWS distinguishes three

multi-agent patterns relevant to finance: **sequential workflows for processes such as AML and claims adjudication, where traceability matters more than speed; swarm patterns for collaborative research across multiple agents;** and graph or hierarchical patterns, which mirror organizational structure and are well suited to loan underwriting, where specialist agents handle credit assessment, fraud detection, and risk modeling under a supervising agent. **Each pattern has its own failure modes. Used in combination, they cover most of the credit value chain.**

What separates the institutions getting durable value from the ones still piloting is rarely the choice of model. It is the discipline with which all three are stitched together inside a single credit operating system, with clear handoffs to human reviewers at every decision point.



Long-term advantage comes from governance, operational discipline, and data quality.

Where the Technology is Showing Up

Across GCC banks, fintechs, and non-bank lenders, the same set of high-impact use cases recurs. Each maps to a specific technical pattern and a specific operational pain point.

- **Origination and Qualification**

Agentic systems are taking over the top of the funnel. They identify prospects from public and licensed data sources, qualify them against a defined credit appetite, gather first-pass documentation through automated outreach, and route only the cases that meet the threshold to human underwriters.

A bank that deploys an agentic origination flow can compress turnaround from weeks to minutes while letting expert underwriters spend their time on the cases that actually need their judgment. One reported deployment processed 175,000 messages and served nearly 5,000 users in six months at a 95% natural language understanding accuracy rate. The same pattern fits retargeting and upsell: agents track the lifecycle, surface refinancing windows, and trigger relationship-managed conversations at the moments where they are most likely to close.

- **Onboarding and Document Processing**

AI-powered integration with third-party data sources, paired with intelligent document processing, replaces what used to be weeks of manual data entry. Financial statements, VAT filings, GOSI records, customs documents, trade invoices, and bank statements are each parsed, classified, and structured in minutes. The lender gets a borrower file that is ready for analysis on day one rather than day thirty. Saudi Arabia's digitized KYC infrastructure, including

national-ID verification and the Chamber of Commerce data fabric, is what makes this loop short enough to matter. Without it, the same workflow takes four times as long.

- **Credit Decisioning**

Retrieval-augmented decisioning sits at the heart of modern underwriting. A RAG layer combines a borrower's structured financials with sector benchmarks, macroeconomic data, regulatory filings, and the lender's own credit policy to produce a draft credit memo that an analyst reviews, refines, and signs off on rather than drafts from scratch. The explainability layer is built in: every claim in the memo cites the document it was drawn from. The human reviewer spends time interrogating the analysis, not assembling it, and the audit trail required by SAMA, the Capital Market Authority, and Shariah boards is generated as a byproduct of the work.

- **Counterparty Screening and Default Early Warning**

Agentic checks run continuously against sanctions lists, politically exposed persons databases, adverse media feeds, and beneficial ownership registries, refreshed in close to real time. PwC and SymphonyAI both report false-positive reductions of up to 80% on entity name screening in this category alone, freeing analysts to focus on the alerts that actually matter.

On the other side of the credit cycle, proprietary machine-learning models read transaction-level behavior for signs of stress: slowing receivable collection, deteriorating debtor mix, atypical drawdown patterns, late salary disbursements, and they flag accounts before the conventional metrics turn. The earlier the warning, the more options the lender has: restructure, top-up collateral, or exit before impairment.

- **The Saudi Institutional Architecture**
The Kingdom's regulators have built one of the most coherent AI-and-data scaffolds in any emerging market. SAMA, the central bank, issued Artificial Intelligence Principles for Financial Institutions in 2023, requiring licensed entities to govern AI use through formal frameworks covering transparency, accountability, fairness, and human oversight. SAMA's Regulatory Sandbox has admitted more than 70 fintechs to date, with over 25 graduating into licensed providers. Its Innovation Hub is itself running generative AI prototypes for supervisory work.
- **SDAIA, PDPL, and Vision 2030**
SDAIA, the Saudi Data and Artificial Intelligence Authority, sits above sectoral regulators and sets national

policy. It launched the National Strategy for Data and AI in 2020, issued two Generative AI Guidelines on January 1, 2024, covering both public and private use, published an AI Adoption Framework with four maturity tiers in September 2024, and put a Deepfakes Guideline alongside it. A Draft Global AI Hub Law went out for public consultation in 2025. The Personal Data Protection Law became fully enforceable on September 14, 2024; SDAIA has issued 48 enforcement decisions in 2025 and 2026 alone, with fines reaching SAR 5 million per breach, doubled for repeat offenses. Sitting above all of this, the Financial Sector Development Program continues to drive the agenda, targeting 525 fintech firms and 18,000 fintech jobs by 2030.



The institutional logic across SAMA, SDAIA, and the FSDP is consistent: encourage adoption, demand governance, and verify through enforcement.

For institutions deploying machine intelligence in credit, that means the regulatory cost is real but the regulatory clarity is rare.

Explainability and the Shariah Question

A model that cannot explain itself will not survive a credit committee in the GCC. SAMA's AI principles, SDAIA's generative AI guidelines, and the PDPL's right to object to fully automated decisions converge on a single expectation: any decision an algorithm informs must be reviewable by a human, and the reasoning must be intelligible to that human.



1 in 3

GCC fintech firms now
use AI in core operations

For complex models, gradient-boosted ensembles, neural networks, deep architectures, accuracy and interpretability often pull in opposite directions. The pragmatic answer is a hybrid stack. The model itself can be complex; an explainability layer (SHAP, LIME, counterfactual explanations) translates each output into the features that drove it. A high-performing model paired with a clear, auditable explanation card is now the working standard in regulated lending. Research consistently finds that **decoupling prediction from explanation** is what unlocks the use of advanced models in compliance-sensitive contexts.

Shariah governance adds another layer. In Islamic finance, contracts must avoid riba and gharar; financing must be linked to real economic assets and identifiable counterparties. AI does not in itself conflict

with these principles. It can in fact strengthen them: automated monitoring can verify that contracts are being deployed correctly, flag exposures to non-compliant activities, and maintain the audit trail Shariah boards expect. But the model has to be auditable end-to-end. A black box that produces a credit decision a Shariah scholar cannot interrogate is unusable, no matter how accurate it is. The institutions that get this right treat explainability not as a constraint on AI, but as a design parameter, built in from the architecture stage rather than bolted on at production. The same is true of fairness testing, model risk management, and ongoing monitoring for drift.

The PDPL, SAMA's principles, and Shariah governance are often described as separate compliance burdens. They are better understood as the same requirement expressed three times: the algorithm must be answerable to the people whose lives and capital it touches. Institutions that internalize this build their AI stack accordingly, with logging, lineage, version control, and reviewable outputs from day one. Institutions that do not will discover the cost the first time a regulator, a Shariah board, or a counterparty asks how a decision was made.



The institutions that turn this expectation into operating discipline early will compound.

The ones that treat it as a future problem will find they have run out of runway by the time the regulator arrives.

The Road Ahead: Controlled Scaling

The trajectory is clear. AI in credit is moving from pilot to production across GCC banks, fintechs, and non-bank lenders. PwC forecasts the regional private credit market expanding at 15 to 30% CAGR over the next five to six years. The infrastructure is being built: Saudi Arabia hosts 33 data centers with 42 more under construction, projected to add roughly 2.2 gigawatts of IT load capacity, nearly a sevenfold increase over current levels. Talent is following the capital, with 86% of Saudi universities now offering AI undergraduate degrees and more than 38,000 students having graduated with AI-related qualifications between 2019 and 2023.

What separates the institutions that compound from the ones that stall is not algorithmic sophistication. The frontier models are commoditizing fast; the underlying capabilities are increasingly available off the shelf. The differentiation lives in the operating discipline around them: clean data, clear ownership, model risk management, governance committees, monitoring for drift and bias, regular Shariah review for Islamic structures, and a culture that treats every model output as a recommendation to a human, not a decision in itself. For investors, lenders, and counterparties, the takeaway is straightforward.

Credit in the GCC is becoming a data discipline. The institutions best positioned to participate in Vision 2030's growth are the ones that combine local credit judgment with the technical capability to read the new signals, and the institutional patience to deploy that capability under controls that regulators, Shariah boards, and capital partners can trust.



Institutions that combine speed, governance, and risk discipline will be best positioned to scale sustainably.

Educational Use & Non-Commercial Disclaimer

This material is provided solely for educational and informational purposes. It has been created to support learning, discussion, and general reference and is not intended to serve as a basis for commercial activity of any kind.

All content presented herein, including but not limited to text, visuals, concepts, designs, and examples, is not offered for sale, resale, licensing, or distribution for profit. Any unauthorized commercial use, reproduction, modification, or redistribution of this material, in whole or in part, is strictly prohibited without prior written consent from the creator or rightful owner.

This material does not constitute professional, legal, financial, or business advice. While every effort has been made to ensure the accuracy and relevance of the information provided, no guarantees are made regarding completeness, reliability, or suitability for any particular purpose. Users are encouraged to seek appropriate professional guidance where necessary.

By accessing or using this material, you acknowledge and agree that it is intended strictly for non-commercial, educational use. The creator assumes no responsibility or liability for any actions taken based on the content provided.

credx

📍 Alfaisaliyah Tower, Office SP02, Riyadh, KSA

✉ info@credx.com

🌐 Credx.com