Design Methodologies for Building AI-native Growth Platforms in Niche Financial Institutions

[1] Pradhan Anubhav
[1] Shiftmate.ai, Co-founder and Chief Product Officer USA

Received: 10th Dec 2025 | Received Revised Version: 19th Jan 2026 | Accepted: 29th Jan 2026 | Published: 16st Feb 2026

Volume 08 Issue 02, 2026 | 10.37547/tajiir/Volume08Issue02-09

Abstract

The article examines design methodologies for constructing AI-native growth platforms in niche financial institutions. The relevance of the study is determined by the rapid transition of financial services toward unified, real-time, and model-centric architectures that overcome the limitations of legacy multi-engine systems. The novelty lies in presenting an integrated analytical synthesis of architectural, operational, and governance principles that jointly define AI-native scalability. The work describes the structural transformation of data pipelines, analyzes constraints of SQL, NoSQL, and NewSQL systems, and studies decisioning, fraud detection, customer intelligence, and regulatory automation workflows. Special attention is given to the role of embedded MLOps and distributed intelligence in sustaining continuous learning. The study aims to identify methodological foundations enabling small institutions to achieve enterprise-grade analytical performance. Comparative analysis, source evaluation, and conceptual generalization are applied. The conclusion outlines an integrated model of AI-native growth and defines practical implications for banks, regulators, and technology designers.

Keywords: AI-native architecture, financial institutions, data pipelines, real-time analytics, decision intelligence, fraud detection.

Cite This Article: Anubhav, P. (2026). Design Methodologies for Building AI-native Growth Platforms in Niche Financial Institutions. The American Journal of Interdisciplinary Innovations and Research, 8(2), 69–77. https://doi.org/10.37547/tajiir/Volume08Issue02-09

Introduction

The accelerating digitalization of financial services has intensified the need for architectural models capable of supporting continuous intelligence, unified data governance, and real-time analytical operations. Niche financial institutions experience these pressures especially acutely, as their historically fragmented data pipelines, batch-oriented processing routines, and isolated analytical engines limit both operational efficiency and competitive responsiveness. The relevance of this research derives from the structural gap between legacy architectures and the requirements of AI-native systems, which depend on persistent data freshness, automated decisioning, and synchronized model governance.

The purpose of the article is to examine how AI-native growth platforms can be systematically designed to overcome the financial, architectural, and operational constraints typical of small and mid-sized institutions. To achieve this purpose, the study formulates three research tasks:

1) to analyze structural limitations of legacy data and analytical architectures in relation to real-time financial workloads;

2) to identify design principles enabling unified ingestion, processing, and decisioning in AI-native platforms;

3) to synthesize architectural, operational, and governance components into a coherent methodological framework applicable to niche institutions.

The novelty of the study lies in integrating heterogeneous strands of research—including data-pipeline engineering, MLOps maturity, distributed intelligence, and financial decision automation—into a unified methodological model. Unlike general surveys of AI adoption in banking, this work focuses specifically on the structural and operational adaptations required for smaller institutions to obtain the analytical and governance capabilities typically associated with large-scale financial organizations.

Methods and Materials

This section summarizes the scholarly and analytical materials used in the study and outlines the methods applied to interpret them. The analysis draws upon peer-reviewed publications and professional reports addressing data-pipeline engineering, MLOps maturity, distributed intelligent systems, and AI-native financial architectures.

The work of Balajee (2025) examines MLOps integration for compliance automation and fraud detection, offering insight into how operational governance shapes model reliability. Burgos et al. (2022) analyze architectural patterns in digital-finance data pipelines, providing detailed descriptions of fragmentation effects in multi-engine environments. Eken et al. (2025) conduct a multivocal review of MLOps practices, identifying persistent challenges in reproducibility and lifecycle management. John et al. (2025) propose a structured maturity model for MLOps adoption, relevant for assessing institutional readiness. Ogenyi et al. (2025) explore distributed intelligence in future connectivity systems, presenting architectural concepts transferable to high-volume financial analytics. Sebastian (2025) investigates resilient insights platforms and explains how real-time analytics transform decision processes in financial services. Somu and Sriram (2023) evaluate next-generation banking infrastructure, offering examples of AI-native IT architectures. Sukharevsky et al. (2025) provide a strategic analysis of agentic AI adoption and its organizational implications.

To address the research goal, the study applies comparative analysis to contrast legacy and AI-native designs; structural analysis to interpret architectural dependencies; and source synthesis to integrate findings into a unified methodological framework. The opening and concluding sections of this study reflect the interpretive generalization of these materials in alignment with the formulated research tasks.

Results

The analysis of design methodologies for AI-native growth platforms in niche financial institutions demonstrates that sustainable performance in constrained banking environments emerges from the convergence of scalable data-pipeline architectures, real-time analytical capabilities, and MLOps-driven operational discipline. These components operate as a single integrated infrastructure in which ingestion throughput, transformation efficiency, and model-governance maturity are inseparable functions supporting credit decisioning, fraud mitigation, customer intelligence, and regulatory compliance. Empirical evidence from recent financial-sector platforms shows that growth in AI-driven services depends on eliminating historical bottlenecks in data movement and enabling architectures capable of processing operational and informational workloads without fragmentation (Burgos et al., 2022).

A key empirical finding concerns the economic and architectural burden of legacy data-pipeline patterns in financial institutions. Organizations with operating expenses between USD 5 billion and USD 10 billion spend USD 90 million–120 million to create and maintain multi-engine architectures due to data fragmentation and duplicated processing layers (Burgos et al., 2022). This cost structure directly affects the feasibility of AI-native transformation in niche institutions, which typically lack the scale to persistently fund such architectures. Below is a structured comparison of legacy and AI-native data-pipeline characteristics (Table 1).

Table 1. Comparative Characteristics of Legacy and AI-Native Data Pipelines (compiled by the author based on Burgos et al., 2022)

Dimension	Legacy Pipeline Characteristics	AI-Native Pipeline Characteristics
Architectural Composition	Multiple disconnected engines; redundant systems	Unified ingestion, processing, and analytics
Data Movement Model	Batch-oriented; high latency	Continuous ingestion; real-time processing
Operational Risk	Frequent inconsistencies; manual reconciliation	Automated lineage and integrity checks
Scalability	Vertical scale limits; fragile under load	Horizontal elasticity; stable ingestion throughput
Fragmentation Impact	Elevated cost, duplicated maintenance	Reduced overhead; centralized governance

Modern reference designs mitigate these constraints by replacing heterogeneous batch-stream pipelines with unified engines supporting real-time ingestion, consistent ACID-level updates, and horizontal scalability without sharding. The capacity to reduce platform costs by up to 30% through architectural simplification, combined with data-infrastructure off-loading and productivity improvements, indicates that AI-native growth strategies rely not merely on new analytical components but on structural redesign of storage and processing layers (Burgos et al., 2022).

A detailed examination of database taxonomies shows that reliance on traditional SQL, NoSQL, and NewSQL systems produces structural tension between ingestion efficiency, transactional integrity, and analytical latency. Niche financial institutions, however, require all three properties simultaneously. The reviewed literature emphasizes that key-value stores offer high ingestion speed but minimal query capabilities, document stores provide schema flexibility but limited ACID consistency, and graph systems become cost-prohibitive at scale due to distributed traversal overhead. NewSQL platforms offer ACID support and improved ingestion but remain constrained under heavy real-time workloads (Burgos et al., 2022). AI-native growth platforms emerge precisely where these data-system properties are fused into unified engines that maintain ingestion efficiency irrespective of dataset size. A representative example is the use of bidimensional partitioning, where historical data are segmented simultaneously by primary key and time dimension so that each partition fits in memory, keeping ingestion time constant and shifting the workload to CPU-bound operations (Burgos et al., 2022). Such architectural shifts allow niche institutions to deploy real-time analytics without replicating the complexity of large-scale banks. The structural distinctions across database paradigms are summarized below (Table 2).

Table 2. Structural Attributes of SQL, NoSQL, NewSQL, and Unified Engine Designs (compiled by the author based on Burgos et al., 2022)

System Type	Strengths	Limitations	Suitability for AI-Native Platforms
SQL	Strong consistency; clear schemas	Limited ingestion throughput	Supports transactional integrity but lacks real-time elasticity
NoSQL	High write speed; flexible models	Weak ACID guarantees; limited complex queries	Useful for ingestion workloads, but insufficient alone
NewSQL	ACID-compliant; improved ingestion	Performance degrades under heavy real-time load	Partial fit; requires architectural augmentation
Unified Engines	Real-time ingestion + ACID + analytics	Requires advanced orchestration	Full alignment with AI-native growth requirements

The literature confirms that real-time decisioning architectures determine growth potential in niche institutions by compressing credit-approval cycles from historical multi-day delays to responses within seconds or minutes (Sebastian, 2025). Embedded decision engines integrate historical bureau data, account histories, bank-transaction patterns, utility-payment behavior, rental records, and in-session behavioral inputs such as form-fill characteristics and device attributes. One extended passage from the reviewed materials highlights that real-time underwriting engines drive buy-now-pay-later offerings and instant loan approvals, reshaping competitive dynamics by drastically reducing latency while guaranteeing transparency and regulatory compliance even during peak transaction volumes (Sebastian, 2025). These engines incorporate ensemble model stacks in which several risk models contribute scored outputs to a final decisioning algorithm calibrated to maintain targeted approval rates and controlled default probabilities. Such architectures, grounded in unified feature stores and automated monitoring pipelines, enable niche institutions to compete with larger players in precision credit delivery.

Fraud-detection pipelines demonstrate parallel benefits. Streaming machine-learning frameworks evaluate millions of transactions per second in sub-second latency windows to identify anomalous activity before settlement (Sebastian, 2025). The literature provides an integrated description showing that fraud systems combine rule-based signals, such as high-risk merchant categories, with deep-learning layers detecting faint anomalies indicative of emerging methods not encoded in explicit rules (Sebastian, 2025). Platforms use continuous monitoring of false-positive and false-negative rates, distribution-shift metrics, and automated retraining triggered when model performance drops below established thresholds. These capabilities, coupled with graph-network analytics for tracing layering schemes and coordinated fraud-ring behavior, indicate that AI-native platforms outperform legacy fraud engines primarily because they bind ingestion, analytics, and retraining into a closed operational loop rather than disconnected batch processes.

Customer-intelligence functions in AI-native growth platforms depend on enterprise-wide unification of transactional, behavioral, interactional, and external datasets via secure APIs (Somu & Sriram, 2023). The reviewed literature shows that integrated profiles combine checking-account activity, credit-card usage, loan data, investment holdings, call-center interactions, branch-office history, and digital-channel behavior, supplemented by credit-bureau updates and life-event signals (Sebastian, 2025). One extended analytic fragment shows that dynamic personalization adapts user interfaces, product suggestions, and educational content to individual behavioral patterns, with machine-learning models refining personalization rules continuously through experimentation and engagement monitoring (Sebastian, 2025). This integrated architecture supports hyper-segmented targeting, propensity-model deployment, retention-risk forecasting, and contextual interventions such as real-time anti-fraud notifications. For niche institutions striving for growth, the convergence of unified customer data and adaptive personalization is a measurable competitive differentiator. The integrated structure of customer-intelligence components is outlined below (Table 3).

Table 3. Components of Customer-Intelligence Architectures in AI-Native Growth Platforms (compiled by the author based on Sebastian, 2025; Somu & Sriram, 2023)

Component	Description	Functional Contribution
Unified Customer Profile	Consolidation of financial, behavioral, and interactional data	Enables holistic visibility across all touchpoints
Behavioral Analytics Layer	Continuous feature extraction from user interactions	Supports adaptive personalization and targeting
Decisioning Engines	ML-driven scoring and recommendation models	Delivers real-time, context-aware interventions
Channel Integration APIs	Standardized access across web, mobile, and assisted channels	Ensures consistent personalization experiences
Monitoring & Feedback Loop	Engagement tracking and rule refinement	Maintains relevance and improves model accuracy

Regulatory-reporting workflows exhibit substantial modernization when incorporated into AI-native growth platforms. Automated lineage capture records every data transformation from ingestion to report generation, replacing manual reconciliation processes that previously consumed extensive analyst time. A detailed passage from the sources notes that major financial institutions achieve quantifiable gains in compliance accuracy and reporting timeliness when regulatory dashboards and audit trails are automated, reducing both cost and operational risk (Sebastian, 2025). These systems maintain data dictionaries, validation rules, exception flows, and multi-framework reporting templates, ensuring that regulatory obligations across jurisdictions are satisfied from a single enterprise dataset. The resulting infrastructure supports operational resilience mandates requiring institutions to maintain key operations, scenario testing, and complete mappings of dependencies across people, processes, facilities, and data assets (Sebastian, 2025). For niche institutions, such architectures reduce the compliance burden through standardization and automated control frameworks.

The reviewed literature on MLOps adoption demonstrates that growth-focused institutions must embed model lifecycle governance directly into platform architecture. The maturity models propose that automated retraining, deployment orchestration, and monitoring pipelines unify experimentation and production environments (John et al., 2025). A comprehensive review identifies unresolved challenges, including environment drift, reproducibility, dependency management, and cross-team coordination, but confirms that integrated MLOps reduces model-degradation lag and improves regulatory transparency (Eken et al., 2025). These findings align with the operational design described in case studies, where continuous monitoring of model drift, real-time retraining triggers, and champion-challenger testing support adaptive credit and fraud systems without manual intervention.

Architectures supporting AI-native operations in niche financial institutions align with broader trends in distributed intelligence. The literature on next-generation connectivity shows how analytic workloads benefit from edge intelligence, semantic data transmission, and real-time distributed decisioning (Ogenyi et al., 2025). While originally oriented toward 6G network environments, these principles demonstrate direct applicability to financial services: local inference reduces central-system load, semantic compression accelerates feature delivery, and distributed decision modules improve latency for fraud detection and credit decisioning in high-volume environments. These patterns reinforce the conclusion that AI-native growth platforms require not only data unity but also distributed analytical execution.

Finally, strategic analyses emphasize that AI-native architectures create measurable economic and organizational benefits. Automation, self-service analytics, and reusable data products reduce infrastructure expenditure, compress time-to-market for new features, and transform fixed costs associated with peak-load management into elastic variable costs aligned with usage (Sukharevsky et al., 2025). Modernization programs that consolidate data stores, retire duplicate systems, and adopt cloud-native designs reduce the total cost of ownership while enabling domain teams to build analytical capabilities independently. This shift allows institutions to redirect capital from infrastructure maintenance toward customer-facing innovation and competitive differentiation (Sukharevsky et al., 2025). These findings confirm that the growth trajectory of niche financial institutions depends on the ability to re-architect data and analytical systems around unified, adaptive, and operationally governed AI infrastructures.

Discussion

The findings demonstrate that AI-native growth platforms in niche financial institutions emerge not from incremental upgrades to legacy architectures but from a structural realignment of data, analytical, and operational layers into a unified system capable of supporting continuous intelligence. The empirical evidence highlights a consistent pattern: the technical and organizational foundations required for sustainable AI-driven growth differ fundamentally from those inherited from batch-oriented, siloed legacy implementations. This divergence explains why smaller financial institutions, despite limited resources, can achieve disproportionate gains when adopting AI-native methodologies, provided that architectural coherence and operational governance are embedded from the outset.

The most significant insight concerns the architectural consequences of fragmentation in traditional data pipelines. Legacy ecosystems rely heavily on multi-engine constructs—combinations of operational databases, data warehouses, data marts, batch ETLs, streaming engines, and intermediate storage—each optimized for a narrow function and poorly integrated with the others. These designs produce latency, inconsistency, and operational risk, particularly when institutions attempt to scale real-time analytics. The Results show that legacy data pipelines not only increase the total cost of ownership but also structurally inhibit the adoption of machine-learning systems that depend on unified data views, low-latency ingestion, and synchronized processing. As a consequence, niche financial institutions often encounter an architectural ceiling: their systems cannot support real-time credit decisioning, fraud detection, or adaptive personalization without disproportionate investment.

The comparative analysis of modern AI-native architectures indicates that the breakthrough lies in the dissolution of the separation between operational and informational data domains. Contemporary platforms unify ingestion, transformation, and analysis into a single substrate capable of sustaining ACID-level transactional consistency while preserving the ingestion throughput typically associated with NoSQL systems. This combination resolves a long-standing tension in financial data systems: the need for strict correctness guarantees and auditability alongside the need for rapid, streaming-intensive analytical workloads. Such architectures do not simply support machine learning; they reorganize the conditions in which machine learning becomes viable.

Real-time credit evaluation, instant underwriting, and dynamic fraud detection depend on this structural realignment, which enables analytical models to operate on fresh, synchronized data rather than stale copies migrated through delayed processes.

A deeper examination of the empirical patterns in credit decisioning and fraud detection reveals that the technical gains of AI-native designs translate into measurable business outcomes. Real-time underwriting pipelines shorten customer-interaction cycles from lengthy manual reviews to decisions delivered in seconds. These improvements are not merely incremental performance gains; they alter product design, competitive positioning, and customer expectations. The ability to integrate heterogeneous data—transaction patterns, payment histories, behavioral indicators, and device attributes—within a single decisioning workflow enables niche institutions to expand access to credit, particularly for thin-file borrowers, without increasing default risk. The presence of ensemble models, continuous monitoring, and automated retraining ensures that decision quality scales with data variety and volume. These advantages are amplified in smaller institutions where traditional underwriting processes are resource-constrained and easily saturated during peak periods.

Fraud detection benefits illustrate a similar transformation. Real-time evaluation of millions of transactions per second is only feasible when ingestion engines, analytical models, and monitoring pipelines function as interconnected components. The shift from rule-based fraud identification to hybrid systems incorporating deep-learning anomaly detection marks a transition from reactive to proactive security. By embedding automated retraining triggers and drift detection mechanisms directly into the platform architecture, institutions gain the ability to adapt to evolving threat patterns without relying on manual updates. This capacity is particularly critical for niche organizations that cannot maintain dedicated cybersecurity research units. AI-native infrastructures effectively amplify institutional capabilities by automating complex investigative and diagnostic functions traditionally performed by specialized analysts.

Customer-intelligence functions demonstrate the broader organizational implications of AI-native architectures. The consolidation of customer data into unified profiles accessible through standardized APIs reshapes marketing, product development, and service delivery. Platforms that integrate transaction histories, interaction logs, life-event indicators, and predictive scores enable hyper-granular segmentation and real-time personalization. This level of precision is unattainable in systems relying on siloed data warehouses or asynchronous batch processes, which limit the frequency and fidelity of customer insights. The dynamic adaptation of user interfaces, recommendations, and learning content in response to behavioral signals illustrates how AI-native platforms convert data into continuous engagement mechanisms that drive loyalty and long-term revenue growth. For niche institutions, where customer intimacy and relationship banking remain competitive differentiators, such architectures accelerate the shift from generic service provision to individualized financial ecosystems.

The transformation of regulatory-reporting workflows further underscores the strategic value of AI-native designs. Automated lineage capture, centralized data dictionaries, and rule-driven exception flows demonstrate how compliance functions evolve when integrated into unified platforms. Instead of manual reconciliation across divergent datasets, institutions can leverage standardized models and automated validation frameworks that adapt to regulatory changes with minimal operational burden. The capability to maintain real-time audit trails, scenario-testing outputs, and cross-framework reporting templates enhances institutional resilience while reducing compliance risk. These gains are particularly relevant for small and mid-sized institutions that must satisfy the same regulatory obligations as large banks but without equivalent resource pools.

MLOps maturity emerges as a fundamental enabler of AI-native growth. The Results confirm that institutions cannot achieve sustainable AI-driven capabilities without automated model lifecycle governance. Reproducibility, dependency management, deployment orchestration, drift monitoring, and retraining cannot remain manual tasks in environments requiring continuous learning. The taxonomy of maturity levels shows that operational excellence requires uniform pipelines enabling reproducible experimentation, controlled deployment, and continuous supervision. By integrating these capabilities directly within the platform rather than distributing them across technical teams, niche institutions reduce operational risk and accelerate the deployment of new analytical services.

The comparative findings on distributed intelligence frameworks indicate that AI-native financial architectures are converging toward hybrid models in which centralized and edge-based analytics coexist. Real-time feature computation at the network edge enhances latency-sensitive functions such as fraud detection and authorization checks. Semantic compression and adaptive data transfer reduce bandwidth consumption and improve the timeliness of analytical inputs. Even though these trends originate from telecommunications research, their relevance extends directly to high-volume financial systems, illustrating the cross-sector convergence shaping next-generation platforms.

Economic and organizational implications reinforce the conclusion that AI-native methodologies provide structural advantages rather than isolated technical improvements. Cloud-native designs transform fixed infrastructure costs into variable operational expenses aligned with demand. Platform consolidation reduces physical footprint, administrative overhead, and technical debt. Reusable data products and feature stores reduce duplicated effort and shorten development cycles. Domain-oriented responsibility models—enabled by unified governance, standardized definitions, and shared data assets—support more rapid experimentation and distribute analytical capability across teams rather than concentrating it in specialized units. For niche institutions with historically limited analytical capacity, this redistribution of capability is one of the most substantial sources of competitive advantage.

The synthesis of these findings suggests that AI-native growth platforms will continue evolving toward deeper integration of decision automation, governance intelligence, and real-time analytical orchestration. Future architectures will likely converge machine-learning workflows, business rules, and human decision pathways into fully unified governance layers. Continuous lineage, real-time compliance monitoring, and embedded explainability will no longer be optional add-ons but inherent components of platform design. The increasing emphasis on fairness, privacy preservation, and risk management indicates that institutional trust will depend not only on model performance but on architectural transparency and operational accountability. Niche financial institutions adopting AI-native designs early will be positioned to overcome the historical disadvantages associated with resource constraints, leveraging platform maturity to achieve agility and innovation historically reserved for large banks.

Conclusion

The study demonstrates that the development of AI-native growth platforms in niche financial institutions requires a fundamental reorganization of data, analytical, and operational structures rather than incremental enhancement of legacy systems. The first research task—identifying structural limitations of traditional architectures—was accomplished through an examination of fragmentation effects, latency constraints, and scalability ceilings inherent in multi-engine designs. These findings clarify why such environments fail to support real-time analytics, continuous learning, or unified governance.

The second task—defining design principles for AI-native systems—was addressed by synthesizing evidence on unified data engines, real-time ingestion models, embedded decisioning workflows, adaptive fraud detection, customer-intelligence architectures, and standardized regulatory automation. The analysis shows that these capabilities converge in platforms built around continuous intelligence and persistent data coherence.

The third task—constructing a methodological framework for niche institutions—was fulfilled by integrating insights from MLOps maturity models, distributed-intelligence research, and strategic modernization studies. The resulting framework illustrates how small institutions can achieve enterprise-grade analytical performance through unified architecture, automated lifecycle governance, elastic scaling, and domain-oriented operational models.

Overall, the conclusions confirm that AI-native design methodologies enable niche institutions to overcome resource constraints, reduce operational friction, and deliver higher-quality financial services. These findings contribute to both academic understanding and practical strategy in the transition toward real-time, model-centric financial infrastructures. The study confirms that AI-native design provides structural advantages beyond technical improvements, offering niche institutions faster credit decisioning, improved fraud detection, personalized customer experiences, and streamlined compliance. Future implementations should focus on phased adoption of unified pipelines, MLOps governance, and distributed intelligence to maximize impact.

References

1. Balajee, A. B. (2025). MLOps in finance: Automating compliance & fraud detection. *International Journal of Computer Trends and Technology, 73*(4), 35–41. https://doi.org/10.14445/22312803/IJCTT-V73I4P105
2. Burgos, D., Kranas, P., Jimenez-Peris, R., & Mahíllo, J. (2022). Architectural patterns for data pipelines in digital finance and insurance applications. In J. Soldatos & D. Kyriazis (Eds.), *Big data and artificial intelligence in digital finance* (pp. 31–52). Springer. https://doi.org/10.1007/978-3-030-94590-9_3
3. Eken, B., Pallewatta, S., Tran, N., Tosun, A., & Babar, M. A. (2025). A multivocal review of MLOps practices, challenges and open issues. *ACM Computing Surveys, 58*(2), Article 39. https://doi.org/10.1145/3747346
4. John, M. M., Holmström Olsson, H., & Bosch, J. (2025). An empirical guide to MLOps adoption: Framework, maturity model and taxonomy. *Information and Software Technology, 107*, 107725. https://doi.org/10.1016/j.infsof.2025.107725
5. Ogenyi, F. C., Ugwu, C. N., & Ugwu, O. P.-C. (2025). A comprehensive review of AI-native 6G: Integrating semantic communications, reconfigurable intelligent surfaces, and edge intelligence for next-generation connectivity. *Frontiers in Communications and Networks, 6*, 1655410. https://doi.org/10.3389/frcmn.2025.1655410
6. Sebastian, D. (2025). Designing resilient insights platforms: Data architecture principles for scalable decision intelligence in financial services. *International Journal of Computational and Experimental Science and Engineering, 11*(4), 7809–7825. https://doi.org/10.22399/ijcesen.4152
7. Somu, B., & Sriram, H. (2023). Next-gen banking infrastructure: Designing AI-native IT architectures for financial institutions. *Journal for ReAttach Therapy and Developmental Diversities, 6*, Article 3610. https://doi.org/10.53555/jrtdd.v6i10s(2).3610
8. Sukharevsky, A., Kerr, D., Hjartar, K., Hämäläinen, L., Bout, S., Di Leo, V., & Dagorret, G. (2025, June). *Seizing the agentic AI advantage: A CEO playbook to solve the gen AI paradox and unlock scalable impact with AI agents*. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage