Cloud-Based Big-Data Architectures for AI-Driven Financial Planning and Analysis and Early Risk Signaling in Resource-Constrained Community Financial Institutions

Naima Bintay Karim[1*], Tomasz Kovac[2]
[1]Arkansas State University, USA
[2]Warsaw Applied AI Laboratory, POLAND

*Corresponding Author Email: naimabintay980@gmail.com

Keywords

Cloud Computing
Big Data
Data Lake
Data Warehouse
Financial Planning and Analysis
Risk Signaling
Community Banks
Credit Unions
CDFI
Cloud Architecture

ABSTRACT

Cloud computing has moved from a peripheral IT decision to the default infrastructure assumption for financial institutions: 91% of financial-services firms globally have embarked on their cloud journey, up from just 37% in 2020, with North America leading at 98% initial adoption, and roughly 90% of financial-services respondents now operate hybrid or multi-cloud environments, against a core banking software market valued at $11.68 billion in 2023 and projected to reach $25.8 billion by 2032 [1][6]. Yet a 2024 cross-industry cloud-maturity survey found only 8% of organisations qualify as highly cloud-mature, indicating a substantial gap between adoption and effective use for most institutions, financial services included [1]. This article examines what cloud-based big-data architecture specifically, data lake, warehouse, and lakehouse patterns can offer resource constrained community banks, credit unions, and CDFIs building AI-driven financial planning and early risk-signaling capability, and what deployment model best fits their specific constraints. The technical case for cloud-based big-data architecture rests on a specific problem: financial institutions generate transaction, core-banking, and increasingly unstructured data (documents, communications, behavioural signals) far too voluminous and too varied in structure for traditional relational databases to handle cost-effectively, yet the risk-signaling and financial-planning use cases addressed elsewhere in this article series depend on exactly this combination of structured and unstructured data [4][5][9]. This article sets out a four-layer architecture ingestion, raw storage, curated storage, and serving evaluates three deployment models against the resource constraints typical of community financial institutions, and addresses the regulatory expectations the FFIEC's Cloud Computing guidance and related examiner practice now apply to cloud-hosted financial workloads.

Introduction

This article draws on cloud-adoption industry surveys, big-data architecture literature, and regulatory guidance on cloud computing in banking to address a practical question for resource-constrained community financial institutions: what cloud-based data architecture actually supports AI-driven financial planning and early risk-signaling capability, and which deployment model fits an institution without a large in-house data-engineering team. Every claim in this article is tied to a specific source rather than presented as general industry consensus.

Cloud adoption in financial services has moved past the experimentation phase documented in earlier industry surveys. Table 1 and Figure 1 set out the scale of this shift: 91% of financial-services firms globally have embarked on their cloud journey (up from 37% in 2020), with North America leading at 98%, and roughly 90% of financial-services respondents report using hybrid or multi-cloud environments,though a

2024 cross-industry cloud-maturity survey found only 8% of organisations qualify as highly cloud-mature [1].

Table 1. Key statistics on cloud adoption and infrastructure scale relevant to community financial institutions.

Statistic	Value	Source
Financial-services firms that have embarked on their cloud journey (global)	91% (98% in North America)	Capgemini World Cloud Report, 2023 [1]
Financial-services respondents using hybrid cloud solutions	90%	Nutanix Enterprise Cloud Index, Financial Services, 2024
Organisations qualifying as highly cloud-mature (cross-industry)	8%	HashiCorp State of Cloud Strategy Survey, 2024
New digital-only banks launching cloud-native	89%	Nutanix Enterprise Cloud Index, Financial Services, 2024
U.S. credit unions (active, 2023)	4,700+, ~$2.3 trillion combined assets	Digital Transformation in Banking market report [7]
Optimized cloud adoption IT cost reduction	Up to ~30%	Zymr core-banking cloud migration analysis [2]

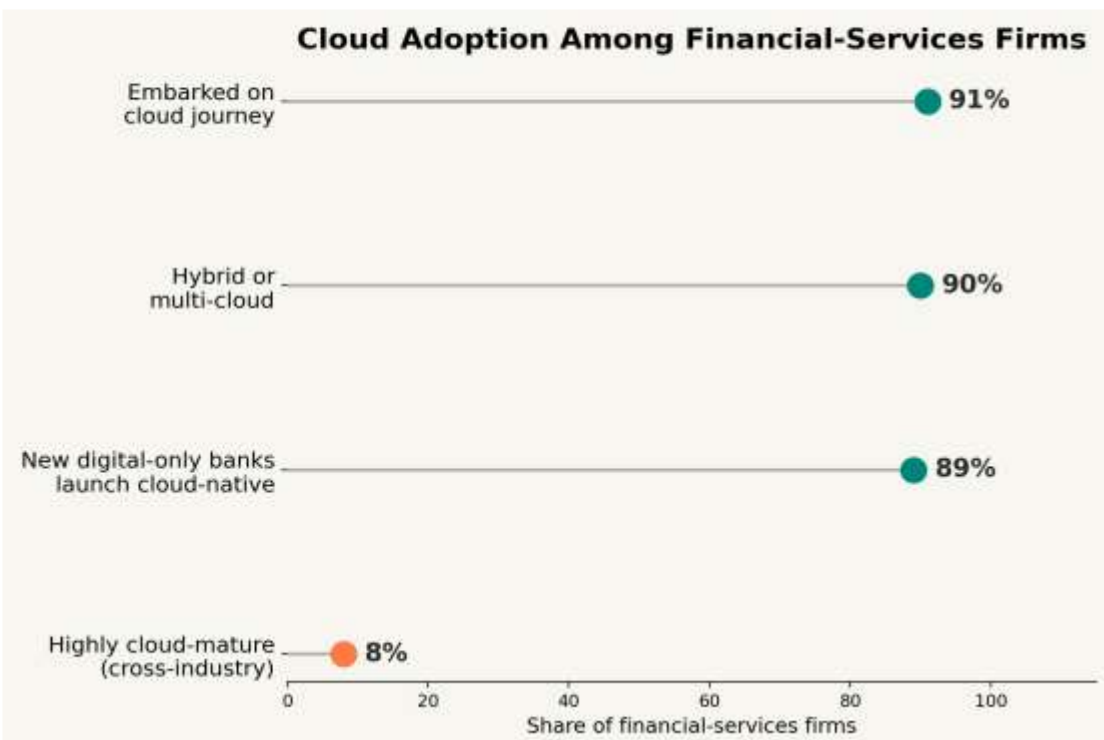

Figure 1. Cloud adoption among financial-services firms [1].

The gap between the roughly 91% adoption figure and the 8% "highly cloud-mature" figure is the central practical problem this article addresses: most institutions have moved some workload to the cloud, but far fewer have built the kind of integrated, big-data-capable architecture that actually supports advanced use cases like AI-driven financial planning and continuous risk signaling, as opposed to simply hosting existing applications on rented infrastructure.

Why Big-Data Architecture, Not Just Cloud Hosting

Migrating an existing application to cloud-hosted infrastructure without redesigning its underlying data architecture captures only a fraction of cloud computing's potential value for AI-driven use cases. The risk-signaling and financial-planning capabilities addressed throughout this article series depend on combining structured transaction data with less structured sources communications, documents, behavioural and transaction-pattern data — that traditional relational databases handle poorly at scale, which is precisely the problem data lake and lakehouse architectures are designed to solve [4][5]. Industry analysis of big data in banking identifies architectural fragmentation multiple legacy systems, siloed data lakes, and vendor-specific platforms producing inconsistent definitions and divergent model outputs — as a primary barrier to realising this value even where cloud infrastructure is already in place [4].

A data lake stores data in its raw, original form, allowing an institution to add new data sources without restructuring existing systems, while a data warehouse or lakehouse layer built on top of that raw storage serves the governed, query-optimised data that reporting and machine-learning models actually consume [9][10][11]. This two-layer pattern is why most mature data organisations use a lake and a warehouse together rather than treating either as a complete solution on its own: the lake provides flexibility and cost-effective raw storage, while the warehouse or lakehouse layer provides the consistency and performance a production risk-signaling model requires [10].

Real-time fraud detection illustrates why this combination matters in practice rather than only in the abstract. A warehouse holding years of historical transaction patterns allows an analytical model to compare a current transaction against an individual customer's typical behaviour and flag a meaningful deviation — the classic example being a debit-card transaction in one city followed by another transaction in a distant city minutes later — almost instantly [9]. This kind of comparison depends on the warehouse layer's structured, query-optimised design; attempting the same real-time comparison directly against a raw data lake, without the curation and indexing a warehouse provides, would be substantially slower and less reliable, which is precisely why the two-layer pattern, not either layer alone, is the architecture this article recommends.

The Data Diversity Problem

The variety of data relevant to modern risk signaling extends well beyond the structured transaction records legacy fraud and credit-risk rules were originally built around. Table 2 illustrates four representative data-source categories a community institution's big-data architecture should be able to accommodate, ranging from the highly structured core transaction records most legacy systems already handle to genuinely unstructured sources most legacy systems cannot process at all.

Table 2. Representative data sources for AI-driven risk signaling, by structure and value.

Data source	Structure	Risk-signaling value
Core transaction records	Highly structured	Baseline behavioural pattern; the foundation most legacy fraud rules already use [9]
Customer communications and service notes	Unstructured text	Early signal of financial distress often expressed before it appears in payment data

News and market-signal feeds	Unstructured text/semi-structured	External context for sector- or region-specific credit risk affecting a business borrower [5]
Scanned KYC and application documents	Unstructured (image/PDF)	Supports document-based verification and fraud checks addressed elsewhere in this series [9]

The unstructured categories in Table 4 are precisely where the data-lake architecture proposed in this article earns its value relative to a traditional relational database. A relational database requires a defined schema before data can be stored, which works well for core transaction records but poorly for a scanned KYC document, a customer service call transcript, or a news feed — data lakes, by storing information in its native format without requiring an upfront schema, allow an institution to begin capturing these sources immediately and defer the question of exactly how to structure and use them until the analytical need is better understood [10][11].

A Four-Layer Reference Architecture

Building on the ingestion-through-governance pattern established in the integration-architecture article in this series, Table 3 sets out a four-layer big-data architecture specifically suited to AI-driven financial planning and risk signaling, extending that earlier architecture's data-ingestion layer into the fuller storage and analytics stack this use case requires.

Table 3. Four-layer cloud-based big-data architecture for AI-driven FP&A and risk signaling.

Architecture layer	Function	Typical technology pattern
Ingestion	Capture transaction, core-banking, and third-party data as it is generated	Streaming pipelines for real-time sources; batch pipelines for periodic exports [8][9]
Raw storage (data lake)	Preserve data in its original form for flexible future use	Object storage holding structured and unstructured data without upfront schema [10][11]
Curated storage (warehouse/lakehouse)	Serve governed, query-optimised data for reporting and modelling	Structured tables built from the raw layer via controlled transformation [9][11]
Serving/signalling	Deliver risk scores and alerts to downstream workflow systems	Low-latency APIs feeding the case-management workflow described elsewhere in this series

The distinction between the raw-storage and curated-storage layers is architecturally important and frequently underappreciated by institutions migrating to the cloud for the first time. Storing everything in a single, undifferentiated cloud database re-creates many of the flexibility and cost problems of a legacy on-premises system, simply hosted on rented infrastructure instead of owned hardware; separating raw, flexible storage from curated, governed storage is what actually allows an institution to add new data sources — a new fraud-signal feed, a new alternative-data source for underwriting without a disruptive schema migration each time [9][10].

Deployment Models for Resource-Constrained Institutions

Table 4 evaluates three deployment models against the specific resource constraints most community banks, credit unions, and CDFIs face: limited in-house data-engineering staff, existing investment in legacy core infrastructure, and a need to demonstrate cloud governance to examiners without a dedicated cloud-compliance function.

Table 4. Cloud deployment models evaluated for resource-constrained community financial institutions.

Deployment model	Description	Best fit for resource-constrained institutions
Full public cloud	All workloads run on a public cloud provider	Newly chartered or digital-only institutions without legacy infrastructure [1]
Hybrid	Sensitive/core workloads stay on-premises or private cloud; analytics run on public cloud	Most community banks and credit unions with existing core investments [3][6]
SaaS point solution	A single vendor-managed cloud service for a specific function (e.g., risk signalling)	Institutions wanting cloud benefits without building or managing infrastructure [6]

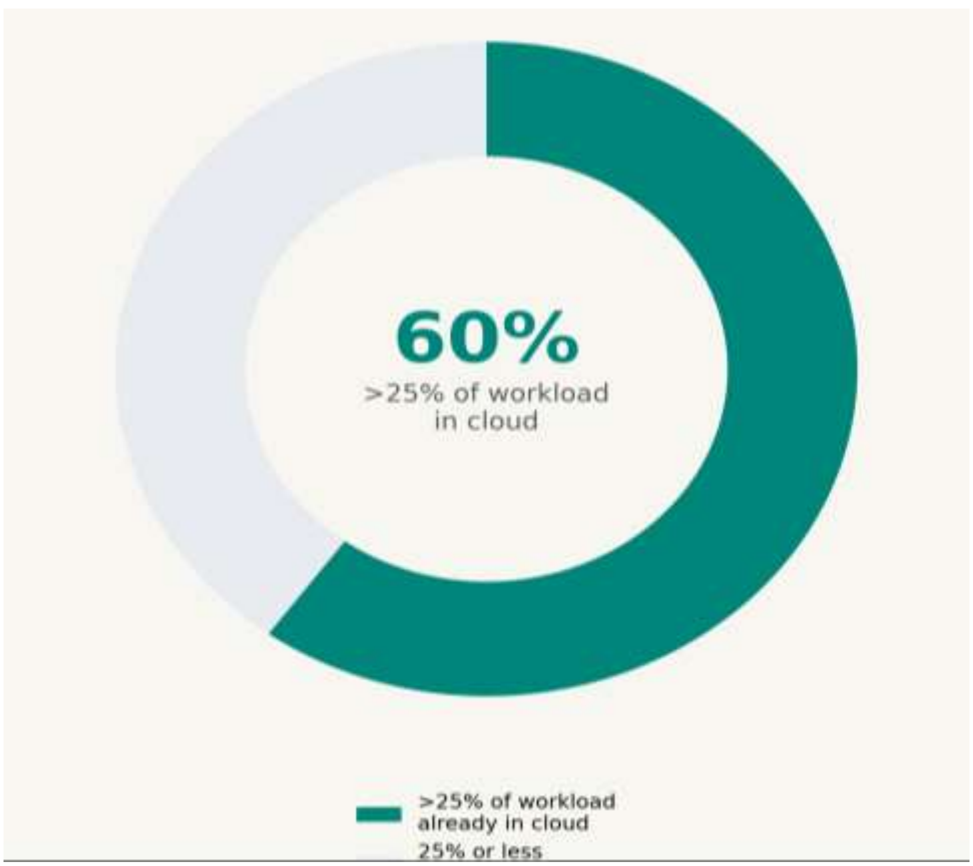

Figure 2. Financial-institution cloud workload adoption, 2023 survey data [2].

For the majority of existing community institutions — as distinct from newly chartered, digital-only institutions that can reasonably build fully cloud-native from inception — the hybrid model remains the most practical starting point, consistent with survey findings that 84% of financial-services executives now recognise cloud and edge computing as a highly relevant trend, with roughly 60% already running more than a quarter of their workloads in the cloud and that share expected to keep rising [2][3]. A hybrid approach allows an institution to keep core transaction processing on infrastructure it already understands and has already satisfied examiners on, while building the big-data architecture in Section 4 on public cloud specifically for the new analytical and risk-signaling workloads this article addresses, rather than requiring the institution to migrate its entire technology estate at once.

The SaaS point-solution model in Table 3 deserves separate emphasis for the most resource-constrained institutions specifically. Rather than building and operating any part of the four-layer architecture in Section 4 directly, an institution can instead adopt a vendor-managed service that delivers a specific risk-signaling or financial-planning capability as a fully managed cloud product, with the vendor responsible for the underlying data architecture entirely. This trades architectural control and customisation for a materially lower operational burden, and is often the most realistic starting point for a CDFI or small credit union without any in-house data-engineering capacity at all — the institution consumes the capability via API or dashboard, addressed in the workflow-integration terms set out in the integration-architecture article in this series, without needing to build or maintain the ingestion, storage, and serving layers itself.

Choosing between the hybrid and SaaS point-solution models is therefore less a technology decision than an organisational-capacity decision. An institution with even a small in-house data or IT function, capable of maintaining the ingestion and storage layers with vendor support, gains more long-term flexibility and lower marginal cost per additional use case from the hybrid model; an institution with no such capacity at all is better served accepting the SaaS model's reduced flexibility in exchange for not needing to build that capacity from nothing. Institutions should make this assessment honestly rather than defaulting to the hybrid model on the assumption that more architectural control is inherently better, since a hybrid architecture an institution cannot adequately staff and maintain is a worse outcome than a well-chosen, properly vendor-managed SaaS alternative.

The global market data in Figure 3 illustrates the trajectory this shift is following: the core banking software market — of which cloud-based platforms are an increasing share — was valued at $11.68 billion in 2023 and is projected to reach $25.8 billion by 2032, a 9.2% compound annual growth rate, indicating that the vendor ecosystem supporting cloud-native and hybrid deployments specifically for banking is still in a relatively early, rapidly maturing stage rather than a settled, mature market [6].

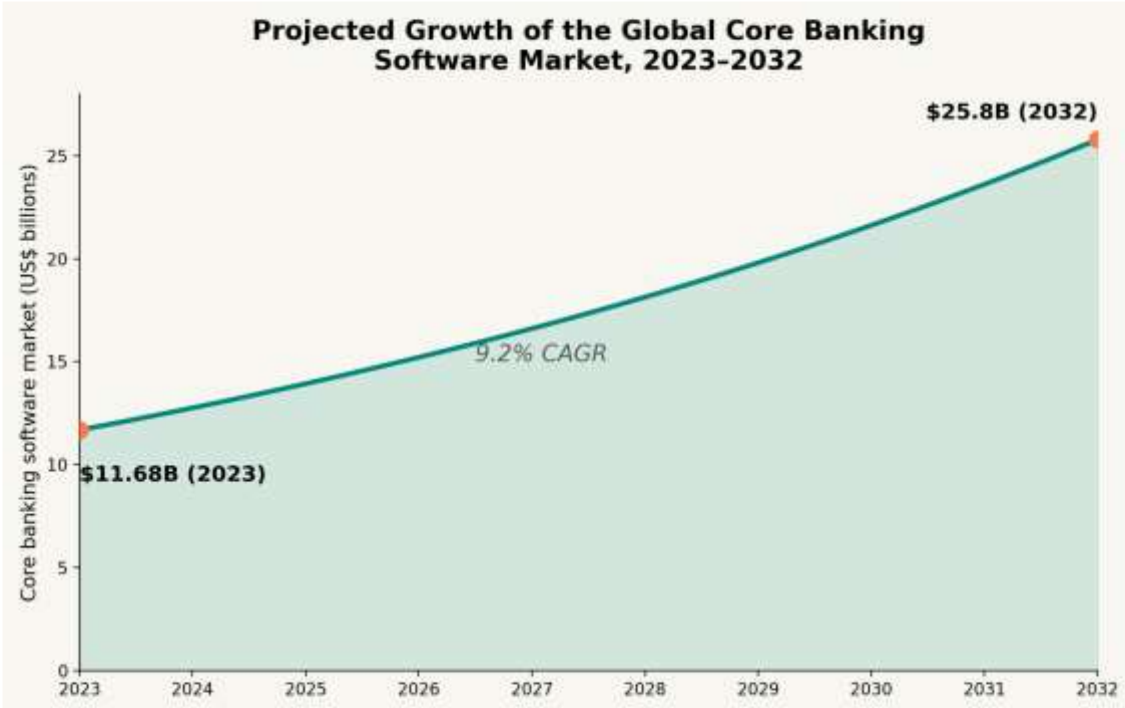

Figure 3. Projected growth of the global core banking software market, 2023–2032 (curve calculated from the cited 9.2% CAGR between the two cited endpoints) [6].

Regulatory Expectations for Cloud-Hosted Financial Data

Moving financial data and risk-signaling infrastructure to the cloud does not reduce an institution's regulatory obligations; it changes how those obligations must be demonstrated. The FFIEC's "Security in a Cloud Computing Environment" guidance makes clear that financial institutions moving workloads to the cloud must demonstrate the same risk-management rigor they apply to on-premises infrastructure, and current examiner practice reflects this directly: NCUA examiners specifically ask credit unions about cloud vendor management and data residency, while OCC guidance and FDIC examination procedures now include dedicated cloud-governance questions [3]. This means an institution cannot treat a cloud migration as purely an IT or vendor-management decision; it is a regulatory-risk decision requiring the same documentation discipline addressed in the model-risk-management article in this series, extended specifically to cover vendor risk, data residency, and third-party access controls.

Reported outcomes from institutions that have already made this transition are notably positive on this specific dimension: IBM's Security X-Force Threat Intelligence Index found that financial institutions using hybrid cloud models experienced 35% fewer security breaches compared to those relying on traditional on-premises infrastructure alone, and HashiCorp's 2024 survey found 66% of organisations increased cloud spending over the past year, up from 56% in 2023, figures shown in Figure 4 [6]. These figures should be read as reported outcomes from institutions that have already navigated the transition successfully, rather than as evidence that cloud migration is inherently lower-risk than remaining on-premises; institutions still early in planning a migration should treat these figures as an achievable outcome contingent on doing the governance work in Section 6 properly, not as a guarantee that adopting cloud infrastructure alone improves compliance and security automatically.

Vendor risk management deserves specific attention within this regulatory discussion because it is the dimension most distinct from an institution's existing on-premises risk-management practice. When core transaction data and risk-signaling infrastructure move to a cloud vendor's platform, that vendor becomes a critical third party whose own security, business continuity, and data-handling practices directly affect the institution's regulatory exposure, in a way a self-hosted on-premises system does not. NCUA and OCC examiners asking specifically about cloud vendor management, as noted above, reflects exactly this shift in risk locus — an institution needs to demonstrate not only that its own architecture is sound, but that it has done adequate due diligence on the specific cloud vendor's practices, has a contractual right to audit or verify those practices, and has a documented contingency plan if the vendor relationship needs to end unexpectedly, none of which has a direct analogue in traditional on-premises infrastructure risk management.

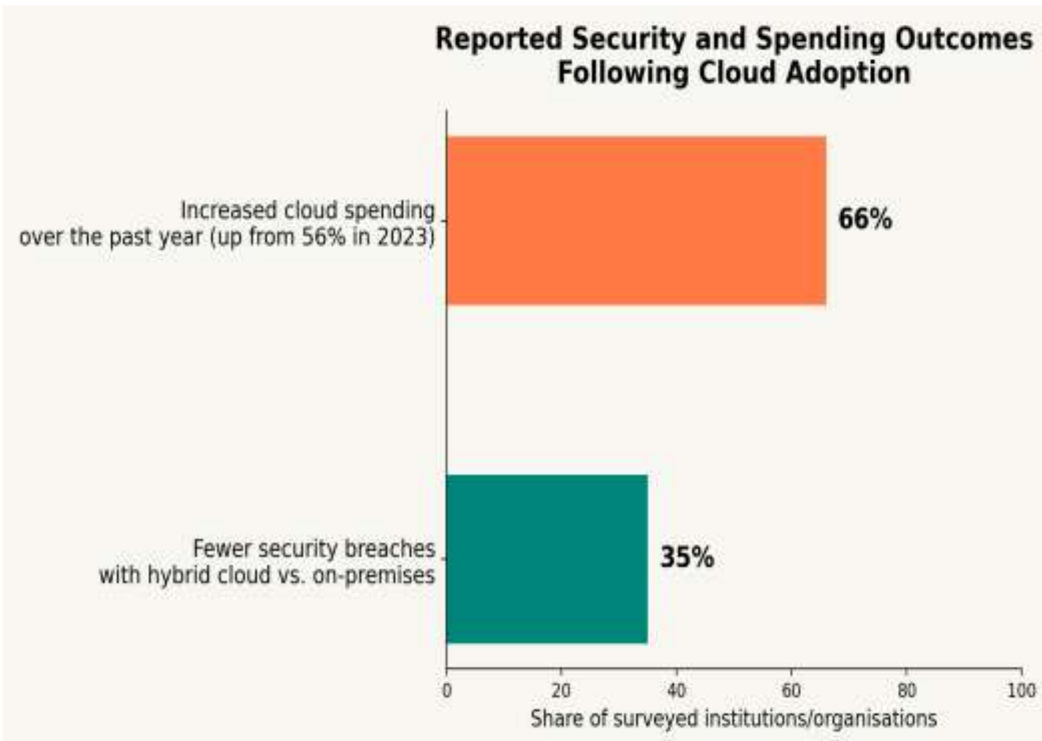

Figure 4. Reported security and spending outcomes following cloud adoption [6].

Cost Considerations and the FinOps Discipline

Cloud infrastructure is frequently, and sometimes inaccurately, assumed to be uniformly cheaper than on-premises infrastructure. Industry analysis is more precise on this point: optimised cloud adoption can reduce IT costs by up to approximately 30%, but poor workload design and the absence of a FinOps (cloud financial operations) discipline can just as easily increase spend relative to an on-premises baseline [2]. This matters specifically for the big-data architecture proposed in Section 4, since a poorly designed raw-storage layer that retains data indefinitely without lifecycle management, or a curated-storage layer that recomputes the same transformations repeatedly rather than caching results appropriately, can generate substantial and avoidable cloud compute costs that a fixed-cost on-premises server would not have incurred.

For a resource-constrained institution without a dedicated cloud cost-management function, the practical implication is that cloud migration planning needs to include cost-monitoring tooling and defined data-lifecycle policies from the outset — archiving or deleting raw data past a defined retention period, rightsizing compute resources to actual query patterns rather than provisioning for worst-case load continuously — rather than treating cost optimisation as a concern to revisit only after costs have already grown unexpectedly large. Given HashiCorp's finding that 66% of organisations increased cloud spending over the past year alone [3], institutions should expect this cost-discipline work to compete for budget against other rising technology costs rather than assuming cloud savings will offset spending growth elsewhere automatically.

A specific FinOps practice worth highlighting for the architecture proposed in Section 4 concerns the raw-storage layer's retention policy. Because a data lake's core value proposition is preserving data in its original form for future, not-yet-defined use, there is a natural temptation to retain everything indefinitely rather than making retention decisions upfront. This temptation should be resisted deliberately: most cloud storage pricing tiers offer substantially lower-cost archival storage for data accessed infrequently, and a defined policy — for example, moving raw data older than 18 months to archival-tier storage while keeping the curated warehouse layer's more compact, actively queried data on standard-tier storage — can capture most of a data lake's long-term flexibility benefit while avoiding the cost growth that comes from treating every byte as equally worth premium-tier storage indefinitely.

A Phased Adoption Path

Given the resource constraints this article addresses throughout, a full simultaneous migration of core systems, data architecture, and analytical capability is rarely realistic for a community institution. A more practical sequence begins with the ingestion and raw-storage layers specifically for the new analytical use cases (risk signaling, financial planning) this article addresses, deliberately leaving core transaction processing on existing infrastructure during this initial phase, consistent with the hybrid deployment model favoured in Section 5. Once this analytical foundation is operating reliably and the institution has direct experience with cloud cost management (Section 7) and cloud-specific regulatory documentation (Section 6), a broader migration — potentially including core systems, following the composable integration architecture addressed in the integration-architecture article in this series becomes a materially lower-risk undertaking than attempting the full scope at once.

A realistic timeline for this phased sequence, based on the deployment-model and architecture considerations discussed throughout this article, typically spans three distinct stages rather than a single migration event. An initial three-to-six-month stage establishes the ingestion and raw-storage layers for one well-defined use case — most commonly the early-warning risk-signaling capability addressed elsewhere in this series, since it has the clearest, most immediately measurable value. A second six-to-twelve-month stage builds out the curated-storage and serving layers for that same use case, incorporating the cost-monitoring and regulatory-documentation practices from Sections 6 and 7 as standard operating practice rather than one-time setup tasks. Only in a third stage, typically beginning after the first use case has operated successfully for at least one full examination cycle, should an institution consider extending the architecture to additional use cases or beginning the broader core-system migration discussed above — sequencing that prioritises proving the architecture works, and that the institution can operate it sustainably, before scaling its scope.

CDFI-Specific Data Architecture Considerations

Community Development Financial Institutions face a distinct version of the resource-constraint problem this article addresses throughout, compounded by the composition issue discussed in the integration-architecture article in this series: a CDFI structured as a loan fund rather than a depository institution may not operate a traditional core banking system at all, and therefore lacks the natural starting point for the ingestion layer in Table 2 that a bank or credit union CDFI would have. For a non-depository CDFI, the practical starting point is instead whatever loan-servicing or portfolio-management system the institution already uses, with the ingestion layer built to extract from that system rather than from a core banking platform.

A further CDFI-specific consideration concerns the mission-outcome data discussed elsewhere in this series jobs supported, borrower demographics, geographic distribution of lending. This data is typically not native to a loan-servicing system at all, often residing instead in separate impact-reporting spreadsheets or a grant-management system used for CDFI Fund compliance reporting. A CDFI's big-data architecture should treat this mission-outcome data as a distinct source feeding the same raw-storage layer as transaction and servicing data, since combining mission-outcome and credit-risk data in a single governed architecture — rather than maintaining them in permanently separate systems is what allows the kind of combined credit-risk-and-mission-preservation monitoring addressed in the fair-lending article in this series to actually be built.

Limitations of the Evidence Base

Several limitations of the evidence assembled here should be stated plainly. Most of the adoption and outcome statistics in Sections 1 through 6 describe financial-services firms broadly rather than community banks, credit unions, and CDFIs specifically, and the largest, best-resourced institutions in any industry survey are likely overrepresented relative to the smaller institutions this article's recommendations are primarily aimed at; a community institution should treat these figures as directional industry context rather than as a benchmark its own smaller program is expected to match immediately. The market-growth projection in Figure 3 extends nearly a decade forward from 2023, and long-range market projections of this kind carry substantial uncertainty that a single cited CAGR figure does not convey. Finally, this article's four-layer architecture in Section 4 is a synthesis of general big-data-architecture literature rather than a proven blueprint validated specifically for the AI-driven risk-signaling use case this article series addresses, and institutions should treat it as a starting framework to adapt rather than a validated, off-the-shelf design.

Recommendations

Four recommendations follow for resource-constrained community financial institutions. First, prioritise building the four-layer architecture in Table 2 for new analytical and risk-signaling use cases specifically, rather than attempting a full core-system cloud migration as a prerequisite, following the phased adoption path in Section 8. Second, default to a hybrid deployment model unless the institution is newly chartered without legacy infrastructure, consistent with current majority practice among established financial institutions. Third, build cloud cost-monitoring and data-lifecycle management into the initial architecture design rather than treating cost discipline as a later optimisation, given the documented risk that poorly designed cloud workloads can cost more, not less, than on-premises alternatives. Fourth, treat FFIEC, NCUA, OCC, and FDIC cloud-specific examination expectations as a first-order design requirement for the architecture in Section 4, integrating cloud vendor-risk and data-residency documentation into the same governance framework addressed in the model-risk-management article in this series, rather than as a separate compliance exercise.

Conclusion

Cloud-based big-data architecture has moved from an optional enhancement to a practical prerequisite for the AI-driven financial planning and risk-signaling capabilities addressed throughout this article series, but the gap between near-universal basic cloud adoption (91% globally, 98% in North America) and cloud maturity (only 8% of organisations qualify as highly cloud-mature) indicates that adoption alone does not guarantee an institution has built the architecture this capability actually requires. For resource-constrained community banks, credit unions, and CDFIs, a phased, hybrid approach building the ingestion, storage, and serving layers in Table 2 for new analytical use cases first, with cost discipline and regulatory documentation designed in from the outset rather than retrofitted later offers a materially lower-risk path to this capability than either remaining entirely on legacy infrastructure or attempting a full, simultaneous cloud migration beyond what the institution's resources can realistically support. Read alongside the other articles in this series, this architecture is the data foundation the model-validation, fraud-detection, integration-architecture, and fair-lending governance work described elsewhere all ultimately depend on: none of those capabilities can operate reliably without the clean, governed, appropriately structured data pipeline this article describes. Institutions planning an AI-driven early-warning or FP&A initiative should treat the data architecture question addressed here as a foundational, sequenced-first decision, not a detail to resolve after model selection and governance design are already underway.

References

[1] Wasserbacher, H., & Spindler, M. (2022). Machine learning for financial forecasting, planning and analysis: Recent developments and pitfalls. *Digital Finance, 4*, 63–88.

[2] Cheng, M., Qu, Y., Jiang, C., & Zhao, C. (2022). Is cloud computing the digital solution to the future of banking? *Journal of Financial Stability, 63,* Article 101073.

[3] Dvorski Lacković, I., Kovšca, V., & Lacković Vincek, Z. (2020). A review of selected aspects of big data usage in banks' risk management. *Journal of Information and Organizational Sciences, 44*(2), 317–330.

[4] Al-Dmour, H., Saad, N., Amin, E. A., Al-Dmour, R., & others. (2023). The influence of the practices of big data analytics applications on bank performance: Field study. *VINE Journal of Information and Knowledge Management Systems, 53*(1), 119–141.

[5] Hussain, A., Nisar, Q. A., Khan, W., Niazi, U. I., & others. (2024). When and how big data analytics and work practices impact on financial performance: An intellectual capital perspective from banking industry. *Kybernetes, 53*(10), 3271–3293.

[6] Lin, E. M. H., Sun, E. W., & Yu, M.-T. (2020). Behavioral data-driven analysis with Bayesian method for risk management of financial services. *International Journal of Production Economics, 228*, Article 107737.

[7] Daníelsson, J., Macrae, R., & Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance, 140*, Article 106290.

[8] Fraisse, H., & Laporte, M. (2022). Return on investment on artificial intelligence: The case of bank capital requirement. *Journal of Banking & Finance, 138*, Article 106401.

[9] Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. *Frontiers in Artificial Intelligence, 5*, Article 779799.

[10] Rupeika-Apoga, R., & Marano, P. (2023). The risk landscape in the digital transformation of finance and insurance. *Risks, 11*(7), Article 129.

[11] Porfírio, J. A., Felício, J. A., & Carrilho, T. (2024). Factors affecting digital transformation in banking. *Journal of Business Research, 171*, Article 114393.

[12] Naimi-Sadigh, A., Asgari, T., & Rabiei, M. (2022). Digital transformation in the value chain disruption of banking services. *Journal of the Knowledge Economy, 13*(2), 1320–1348.

[13] Jihane, T., & Aziz, M. (2022). Banks and FinTech relationship in a digital transformation context. *European Scientific Journal*, *18*(12), 106–122.

[14] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLOS ONE, 13*(3), Article e0194889.

[15] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting, 36*(1), 54–74.

[16] Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review. *Applied Soft Computing, 90,* Article 106181.

[17] Hewamalage, H., Bergmeir, C., & Bandara, K. (2021). Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting, 37*(1), 388–427.

[18] Bandara, K., Bergmeir, C., & Smyl, S. (2020). Forecasting across time series databases using recurrent neural networks on groups of similar series: A clustering approach. *Expert Systems with Applications, 140*, Article 112896.

[19] Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing, 50*, 159–175.

[20] Cheng, X., Zhang, Y., & others. (2024). Data privacy and cybersecurity challenges in the digital transformation of the banking sector. *Computers & Security, 147*, Article 104051