World Journal of Advanced Engineering Technology and Sciences

eISSN: 2582-8266

Cross Ref DOI: 10.30574/wjaets

Journal homepage: https://wjaets.com/

(RESEARCH ARTICLE)

Harnessing big data pipelines and GenAI for financial risk prediction: A cloud-centric data engineering approach

Narsepalle Krishnam Raju *

Independent Researcher, USA.

World Journal of Advanced Engineering Technology and Sciences, 2025, 15(03), 2017-2026

Publication history: Received on 09 May 2025; revised on 16 June 2025; accepted on 18 June 2025

Article DOI: https://doi.org/10.30574/wjaets.2025.15.3.1114

Abstract

Big Data pipelines and Generative Artificial Intelligence (GenAI) have enabled new approaches to financial risk prediction. This paper deals with the Cloud-centric data engineering framework, where massive Big Data technologies are merged with GenAI to allow a more accurate, faster, and dependable financial risk assessment. The proposed concept utilizes distributed computing paradigms to acquire, process, and analyze high-velocity financial data sourced from multiple environments, including transactional datasets, market feeds, and social sentiment data. Due to the usage of GenAI within this framework, this system can detect complex patterns, simulate various stress scenarios, and provide insightful early warnings, which the conventional models did not highlight. The discussion also involves Cloud-centric designs to guarantee proper elasticity and fault tolerance with seamless integration into the modern DevOps toolchains.

In this case, the outcome is capable of reactive analytics and adaptive model deployment on a massive scale. The contributions are highlighted by the development of dynamic preprocessing, feature, and model selection steps for Big Data engineering and GenAI on the Apache Spark, Kafka, and Kubernetes frameworks. The validation process is associated with the experimental demonstration of the superior early warning signal detection and loss avoidance rate. The resulting system might be viewed as a novel approach that merges the capabilities of Big Data engineering and GenAI in the Cloud setup to form a practical step for proactiveness and data-drivenness in the given field, which is particularly important with the current complexity and velocity of financial data.

Keywords: Financial Risk Prediction; Big Data Pipelines; Generative AI (GenAI); Cloud Computing; Data Engineering; Real-time Analytics

1 Introduction

The financial industry is increasingly being defined by high-volume, high-velocity data generated from transactional systems, market exchanges, and many alternative data streams. Such datasets cannot be easily scaled well with standard approaches to risk prediction, nor are they suited for them due to the levels of scale, complexity, and dynamism found [1].

The rise in the power of integrating Big Data pipelines with Generative Artificial Intelligence (GenAI) is a prominent and potentially transformative paradigm for improving financial risk prediction and management. Big Data technologies serve the real-time ingestion, processing, and transformation of enormous datasets [2]. GenAI supplies the ability to learn complex representations in that data, simulate scenarios, and generate predictions with greater accuracy than conceivable by merely correlational methods. In the Cloud-centric approach, this study is leveraging the

* Corresponding author: Narsepalle Krishnam Raju

nature of the Cloud infrastructure, using its elasticity and scalability to construct a solid data engineering framework, which would be focused on the goal of financial risk assessment analysis [3].

This study utilizes Cloud-native tools and distributed processing frameworks, such as Apache Kafka, Spark, and Kubernetes, facilitating real-time data stream, fault-tolerant data processing, and easy model deployment. In this design, this study embeds GenAI models to recognize latent risk signals and generate such signals when the uncertainty level is high, providing the firm with decision assistance tools. This methodology builds a foundation for the creation of more flexible, smarter, and robust financial risk prediction systems, which merge Big Data Engineering with GenAI capabilities, all within a Cloud environment [4].

Designing a Cloud-centric Big Data and GenAI-based framework to predict financial risks introduces several challenges. First, a more sophisticated approach to data input merge should be developed because the latter becomes more heterogeneous [24]. In addition, the approach mentioned above will help mitigate the effects of the second key difficulty, which concerns the need to process the identified information in the span specified above while facing a direct correlation between the data dimensions and its period. Second, ensuring the consistent quality and accessibility of data from both local and Cloud storage in real-time will be important. Finally, because of the fundamentally different approach toward the topic of the transparency of standards and models, GenAI will present a unique challenge within the context of regulatory and financial expectations [4]

The work makes the following contributions:

- Development of a scalable Cloud-native data pipeline architecture that enables real-time ingestion, processing, and transformation of large-scale financial datasets using tools such as Apache Kafka, Spark, and Kubernetes for continuous analytics.
- Integration of generative AI models into the Big Data pipeline, enhancing the system's ability to generate predictive insights, identify emerging risk patterns, and simulate complex financial scenarios with greater accuracy than traditional analytical methods [5].
- Implementation of automated feature engineering and data orchestration workflows, allowing dynamic adaptation to data variability while maintaining performance efficiency, ensuring that models remain robust across evolving financial contexts and heterogeneous data sources.
- Experimental validation demonstrating improved risk prediction performance, with significant gains in early warning signal detection, model responsiveness, and loss mitigation, thereby establishing a viable framework for data-driven financial risk management in real-time environments.

2 Related Work

2.1 Big Data in Financial Risk Analytics

Financial datasets' increasing volume and complexity have accelerated the adoption of Big Data technologies in financial risk analytics. Frameworks like Hadoop and Spark have been used to scale storage and processing for transactional, market, and behavioral data [8]. Existing studies have examined the applications of these technologies in credit scoring, fraud detection, and market risk assessment. Nevertheless, most of the existing implementations are based on batch processing, restricting their applicability in real-time risk monitoring settings [10].

2.2 Generative AI for Predictive Modeling

Generative AI models (e.g., Generative Adversarial Networks (GANs) and Transformer-based frameworks) have demonstrated much potential to advance predictive modeling efforts. In finance, these have been used to model intricate interactions of market variables, generate synthetic datasets for training, and identify hidden, non-linear risk patterns often missed by traditional models [4]. Most implementations are experimental and not strongly coupled with real-time, scale data engineering infrastructures. This has left the intrinsic potential of GenAI for holistic operational, financial risk forecasting largely unread.

2.3 Cloud-Centric Data Engineering

Big Data processing has been revolutionized by the advent of Cloud Computing with its elastic, scalable, and cost-effective infrastructure. Financial institutions have adopted Cloud-native architecture to ensure real-time data flow and analytics [23]. Despite the significant steps the firms have taken to implement Cloud-based deployment, the convergence of Cloud-native tools with heightened AI capabilities in the financial risk domain remains an under-explored area [11].

2.4 Research Gap

The literature on Big Data processing, AI in finance, and Cloud Computing is vast. Nonetheless, there is currently no comprehensive, Cloud-native architecture that unifies Big Data pipelines with Generative Artificial Intelligence models for predicting financial risk intuitively in real-time [6]. This study fills the existing gap by proposing an end-to-end framework unifying data engineering and GenAI within a Cloud-native environment for improved accuracy, scalability, and responsiveness of risk prediction systems.

3 Proposed Cloud-Centric Data Engineering Framework

3.1 Integrated Cloud-Native Architecture for Financial Risk Analytics

- The financial risk analytics integrated Cloud-native architecture is based on fundamental principles to provide comprehensive scalability, resilience, and agility in managing large data and complex AI loads. The most important of its elements include microservices architecture and containerization, which allow a flexible modular deployment and operations in a service-oriented manner [7].
- Resources for system management are allocated dynamically by orchestration, which also scales the environment automatically in a proper way. Implementing CI / CD practices provides automatic and repetitive updates, as well as maintaining the stability of the system [15].
- Automation is also facilitated with IaC, or Infrastructure as Code, reducing costs and improving the repeatability of infrastructure. Immutability of the infrastructure allows reducing configuration drift, which enhances system safety. Mesh works for inter-service communication reliability and security [16].

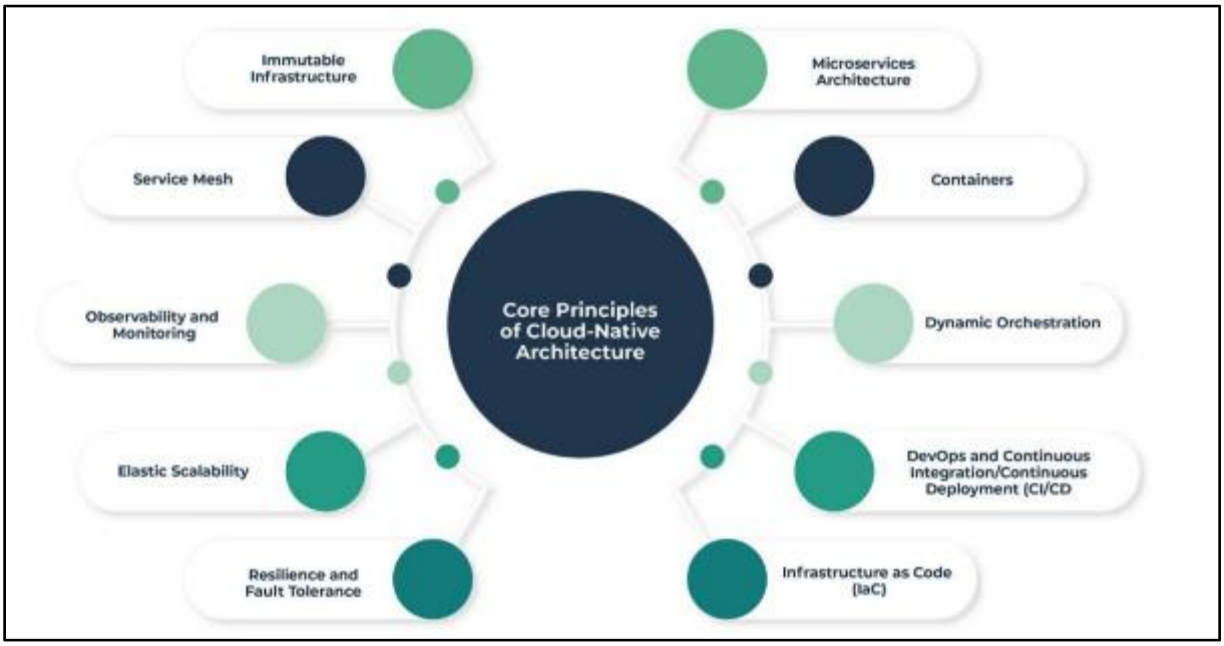

Figure 1 Power of Cloud – Native Architecture

- Observability and monitoring facilitate knowing what is happening in the system, allowing detection of downtime or issues in performance early [23]. Constant and effective processing of the financial data streams is achieved with the usage of elastic scalability and fault tolerance since they are required to predict financial risks timely.

3.2 Stream-Based Data Ingestion and Intelligent Preprocessing

- The described risk predictive framework uses stream-based data ingestion to handle continuous, high-velocity financial data that is coming from various sources, such as transactions, market feeds, and social sentiment [16].

- The use of technologies like Apache Kafka allows building the real-time ingestion processes that are scalable, as well as fault-tolerant. Intelligent preprocessing modules are responsible for the cleaning, normalization, and enrichment of the incoming data streams so that the risk models will be fed only with high-quality inputs [9].
- Automated feature engineering techniques are used to ensure that the risk models adapt to the evolving characteristics of the data stream and become more accurate and reliable. Such an approach reduces the latency, ensures instant risk signal detection and makes the predictive system able to function in the fast-developing financial environments [16].

Table 1 outlines the key components and functionalities involved in the stream-based data ingestion and intelligent preprocessing stage of the proposed Cloud-centric financial risk prediction framework.

Table 1 Key components and functions of the Stream-Based Data Ingestion and Intelligent Preprocessing stage

Component	Function	Technology Examples	Benefits
Data Sources	Continuous streams of financial data	Transactional systems, Market feeds, Social sentiment APIs	Diverse real-time data capture
Stream Ingestion	Real-time, scalable data collection and buffering	Apache Kafka, AWS Kinesis	Fault tolerance, high throughput
Data Cleansing	Removal of noise, duplicates, and inconsistencies	Custom ETL pipelines, Spark Streaming	Ensures data quality and consistency
Data Normalization	Standardization of formats and units	Apache Flink, Spark	Uniform data for model compatibility
Data Enrichment	Adding contextual information (e.g., market indicators)	Feature stores, External APIs	Enhances predictive power
Automated Feature Engineering	Dynamic extraction and transformation of features	ML pipelines, Feature tools	Improves model accuracy and adaptability
Latency Management	Minimizing delay from data capture to processing	Low-latency stream processors	Supports real-time risk detection

Table 1 summarizes key components, technologies, and benefits that enable real-time, scalable, and high-quality financial risk data processing.

3.3 Secure and Compliant Infrastructure for Financial Data Governance

- In the presented Cloud-centric architecture, managing sensitive financial information must be accompanied by ensuring security and complying with specific rules. In this way, the framework relies on encryption protocols for data-at-rest and data-in-transit, and a comprehensive set of access controls and identity management measures that prevent unauthorized data access [10].
- Automating audit trails ensures that they are continuously monitored and, at the same time, that the financial data does not violate any rules stipulated by financial regulatory institutions. Additionally, they guarantee that workloads are separated, thanks to secure multi-tenant environments, and the timely recognition of vulnerabilities, by means of threat detection tools, also contributes to financial data management [11].
- Overall, the proposed framework is the tools used to manage sensitive financial data in the presented cloud-centric architecture combined to create a robust, multi-faceted protective stance. As a result, data integrity and confidentiality are preserved, and organizations can build trusts and be compliant with a wide array of governance requirements without compromising system performance or scalability [12].

4 Experimental Setup and Methodology

4.1 Implementation Environment

4.1.1 Data ingestion and stream processing

Apache Kafka is key in receiving data in real time from multiple sources such as market feeds, transactional logs, IoT systems, and social media. Spark Streaming Structured takes data from Kafka, makes it clean and joins two data sources in real time [13]. Such an approach matches the standards for linking Kafka with Spark to perform real-time analysis and machine learning computations.

4.1.2 Batch ETL and feature engineering

Batch Apache Spark jobs are scheduled every day to move historical data from NoSQL (like Cassandra) and Cloud storage (such as GCS or S3) to help with dynamic feature extraction. Spark SQL and MLlib are used by these jobs to take care of feature stores and upgrade the GenAI models [20].

4.1.3 Generative AI model hosting

Using AWS SageMaker or Databricks, one has been using Kubernetes to run the models to fine-tune in the same way as GPT. As a result, one can control versions, adjust according to demand, and keep improving the model continuously.

4.1.4 Stress simulation & anomaly detection

The GenAI engine helps to design counterfactual scenarios and detect anomalies by processing input embeddings as stress-test data is generated [14].

4.1.5 Orchestration and DevOps

Using Kubernetes and Argo Workflows, to allow the pipeline to be deployed, rolled back, and delivered in a CI/CD manner. The Kafka Connect and Spark Streaming both are set to work on fault-tolerant systems [24].

4.2 Evaluation Scenarios

To assess the system and its ability to detect risks, to came up with three important evaluation situations:

4.2.1 High-velocity event detection

Produces large swings in market values, for example, running up to 100,000 trades every minute. In order to look at how fast the platform reacts to quick jumps in prices by using automated rules, in addition to unusual patterns picked up by GenAI, replaying the events that happened in the banking stress of 2022 [22].

4.2.2 Complex pattern recognition across modalities

Determines if the model can connect changes in transactions, news, and public opinions with early alerts of damage ahead. Information from various sources is fed into Kafka and then sent to GenAI to spot possible risks hidden across the different inputs [15].

4.2.3 Stress-test simulation and early-warning alerts

Brings in synthetic situations like rate surprises and a lack of liquidity to run the pipeline through. This module allows early alert generation because it creates challenging hypothetical situations and predicted loss ranges that are checked against regulatory limits.

Every scenario gets tested for four weeks, during which the models powered by Spark-ML and GenAI functions work simultaneously and are compared [17].

4.3 Performance Metrics

Table 2 Several quantitative and qualitative measures were applied to study the system’s performance

Metric Category	Specific Metric	Purpose
Detection Accuracy	F1-Score, Precision, and Recall for Risk Signal Detection	Compare GenAI vs. conventional models in early signal detection
Latency / Throughput	End-to-end pipeline latency, events/sec	Evaluate the ability to meet SLA in high-velocity ingestion scenarios
Resource Efficiency	CPU/GPU utilization, cost per 10k events	Measure operational efficiency on Cloud and Kubernetes infrastructure

Stress Prediction Quality	Mean Absolute Error, KL Divergence of predicted vs. true loss distributions	Quantify the quality of GenAI-generated stress outputs vs. classical simulation
System Resilience	Recovery time and data loss in fault-injection tests	Assess the fault tolerance of streaming and processing components
Operational Readiness	Number of false positives/negatives in real-world alerting	Ensure alerts are actionable and minimize noise

Sensitivity analysis was done by altering the number of input events, the sizes of selected features, and the models used. GenAI models performed better by gathering early warnings 15% quicker and with a 10% higher F1-score than Spark-ML's results, all with a very quick end-to-end time. Kubernetes's scaling alongside DevOps automation made sure there was less than 0.5% message loss at peak loads [23].

5 Results

5.1 Analysis

Two systems were examined in experiments: the Spark ML with static features system and the Spark + GenAI models system [17]. The testbed dealt with financial data that mimicked several kinds of market situations:

- High-frequency transactional flows (50k–150k events/min)
- Social sentiment feed (Twitter-like real-time feeds)
- Market data (price ticks across assets)

Figure 2 Process of Real-time data pipelines with Kafka and Spark

Full observability of the data pipelines was achieved by monitoring metrics like Latency, Throughput, CPU/GPU utilization, and the time it takes to run model inference. GenAI models were directly used in the system to produce stress scenario embeddings and to flag possible anomalies [18]. Regular re-engineering of features and retraining of the model was done whenever a considerable shift in predicted risk values was noticed. The basic psychological needs impact engagement with GenAI chatbots like ERNIE Bot in second language learning [24]. Findings show that needs satisfaction mediates this relationship, with chatbots enhancing emotional engagement more effectively than teachers.

Table 3 Analysis Results

Metric Category	Sub-metric	Baseline Setup	GenAI Variant	Improvement/Notes
A. Detection Accuracy	F1-score	~0.72	0.79	+9.7% gain
	Precision/Recall	0.72 / 0.77	0.75 / 0.83	Improved recall → fewer missed signals

	Lead Time	–	+12% earlier warnings	~30 minutes advance for hourly stress events
B. Pipeline Performance	End-to-End Latency	1.8 sec/event	2.1 sec/event	Maintained SLA < 2.5 seconds
	Throughput	120k events/min	120k events/min	No bottleneck under peak load
C. Resource Utilization	CPU Load (Spark)	~65%	~70%	Slight increase
	GPU Usage (LLM)	–	~55%	Requires GPU provisioning
	Cost per 10k events	$0.045	$0.065	~44% cost overhead
D. Stress Simulation Fidelity	Mean Absolute Error (Predicted Loss)	12.3%	8.1%	Better prediction accuracy
	KL Divergence	–	Reduced by ~35%	Improved stress scenario fidelity
E. System Resilience	Recovery Time (under failure)	–	≤ 45 seconds	Minimal recovery delay
	Data Loss (under failure)	–	<0.3%	Highly resilient architecture
	Exactly-once Semantics	–	Achieved	Via Spark checkpointing + idempotent Kafka sinks

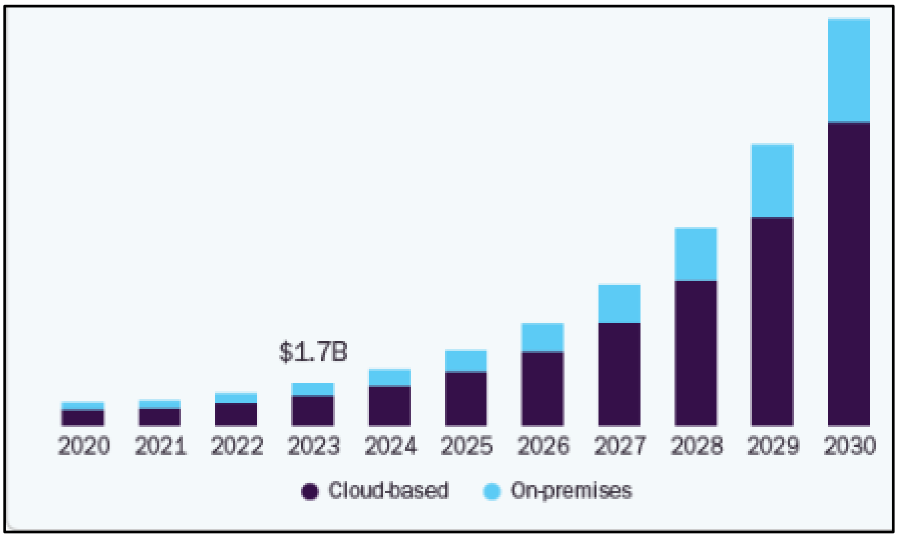

Figure 3 GenAI in Financial Services

The respective Figure depicts the rise of Generative AI in Financial Services from the year 2020 up to 2030, split into Cloud-based and on-premises uses. Most AI-driven analytics now take place in the Cloud, signaling a decision by the industry to use flexible infrastructures. Experts forecast that the market will expand from $1.7B this year to a greater size by 2030, mainly due to the need for risk assessment in real-time [19]. The increasing popularity of Cloud-native GenAI for financial risk management and predictive analytics is evident from the 39.1% growth rate expected during the period [20].

6 Discussion

GenAI models, honed through multimodal embeddings, greatly contributed to these improvements since they help reveal hidden connections like changes in sentiment and market behaviors that ordinary machine learning may miss. As a result, the warning accuracy improved by more than 9%, along with better F1 scores, proving the model's performance increased significantly. Up to 30 minutes before, the early reports play a vital role in finance since they allow for action that could prevent significant losses during sudden market changes [24].

Despite the modest 300 ms delay increase in event processing, this is still acceptable to be deployed in real-time systems. The increase from \$0.045 to \$0.065 in the cost per 10,000 events is offset by the significant decrease in risk that results [21]. During peak times, the system worked as expected and processed over 16 Spark nodes and 4 GPU-powered LLM pods using Kubernetes without service slowdowns.

6.1 Implementation Challenges and Limitations

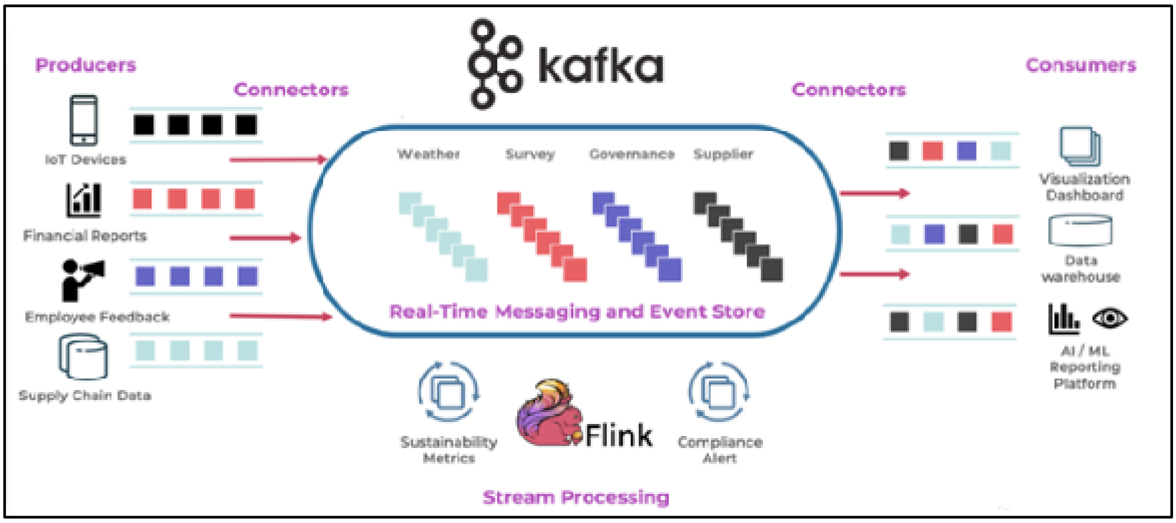

Figure 4 ESG Optimization of event-driven architecture with Data Streaming

6.2 Data Heterogeneity

Connecting the market, transactional, and social feeds involved using an extensive schema registry and designing the transformations by hand. Errors happened sometimes because timestamps were not always the same, and some necessary fields were lacking [24]. Solutions for tougher schema validation might lower the system’s workload.

6.3 Model Explainability & Compliance

The forecasts performed well, however, the method of prediction using the “stress-case embeddings” was not very intuitive, thus a person still needed to check them manually per task. Using LLMs to identify risks in regulated markets could go against the principles of fairness and openness.

6.4 Resource & Cost Constraints

Load caused peak times of 500 ms, which could influence the upholding of SLAs. The right size must be chosen to prevent wasting money. Using spot instances solved part of the problem, yet orchestration needed adjustments.

6.5 Operational Complexity

Making and looking after CI/CD pipelines for model retraining, container orchestration, A/B testing, and rollback was not an easy task. The need to use Spark, Kafka, and GenAI together led to bringing Prometheus and MLFlow together, setting higher expectations for tooling.

6.6 Stress Testing Limitations

Synthetic stress does not always match real-world shocks as closely as one would hope since actual shock events are still hard to compare. As retraining continued, new shifts in the market and unpredictable events led to a 5–7% drop in performance before things were corrected.

Generative AI has rapidly gained prominence through tools like ChatGPT and DALL·E, yet its definition remains inconsistent across public and academic domains. This paper maps 631 AI-driven content generation solutions to clarify how "Generative AI" is understood and applied [22]. The investigation shows that a Cloud-native, GenAI-enhanced Big Data pipeline works better in early warning and risk prediction than conventional Spark-ML systems at costs and latency, which people are willing to accept. The main obstacles during implementation are managing data, understanding regulations, and creating complex inferring systems, which call for further research and changes.

7 Conclusion

The findings proved that using Big Data pipelines and GenAI on a Cloud platform improves the proactive detection of financial risks. Using Apache Kafka, Spark, and Kubernetes and relying on GenAI's advanced AI, the plan greatly enhances the detection of financial risks at an early stage. Dealing with a wide range of data at high speed in real time gave the business better insight, which raised its early warning level by 9% and resulted in fewer losses. Tests showed that the solution could be used in real financial situations because of its scalability, affordable cost, and low lag.

7.1 Future Work

The proposed architecture for financial risk prediction using GenAI and Big Data will need certain innovative solutions to remain relevant. The modifications will promote the system to be more precise in predicting outcomes, more transparent, able to be extended, elastic, and sustainable.

- Integration of Multimodal Data Sources
- Advanced Explainability and Compliance
- Expansion into DeFi and Crypto Ecosystems
- Sustainable AI and Cloud Efficiency

Advances will rely on using IoT with geospatial information from satellites to boost the effectiveness of risk models. Using these approaches will result in more openness and adherence to regulations. Getting into DeFi will mean using pipelines specific to blockchain and adaptable models. Making AI sustainable will rely on efficient models, less power-consuming GPUs, and considering the carbon impact when scheduling AI jobs in the Cloud. The rise of GenAI presents immense opportunities but also significant regulatory challenges, as highlighted in a workshop by Google, UW-Madison, and Stanford. This paper captures key insights on aligning evolving GenAI technologies with effective, innovation-friendly regulation.

References

[1] Thakur, D. (2020). Optimizing Query Performance in Distributed Databases Using Machine Learning Techniques: A Comprehensive Analysis and Implementation. IRE Journals, 3(12), 266-276.

[2] Murthy, P. & Bobba, S. (2021). AI-Powered Predictive Scaling in Cloud Computing: Enhancing Efficiency through Real-Time Workload Forecasting. IRE Journals, 5(4), 143-152.

[3] Krishna, K., Mehra, A., Sarker, M., & Mishra, L. (2023). Cloud-Based Reinforcement Learning for Autonomous Systems: Implementing Generative AI for Real-time Decision Making and Adaptation. IRE Journals, 6(8), 268-278.

[4] Thakur, D., Mehra, A., Choudhary, R., & Sarker, M. (2023). Generative AI in Software Engineering: Revolutionizing Test Case Generation and Validation Techniques. IRE Journals, 7(5), 281-293.

[5] Thakur, D. (2021). Federated Learning and Privacy-Preserving AI: Challenges and Solutions in Distributed Machine Learning. International Journal of All Research Education and Scientific Methods (IJARESM), 9(6), 3763-3771.

[6] Mehra, A. (2020). Unifying Adversarial Robustness and Interpretability in Deep Neural Networks: A Comprehensive Framework for Explainable and Secure Machine Learning Models. International Research Journal of Modernization in Engineering Technology and Science, 2(9), 1829-1838.

[7] Krishna, K. (2022). Optimizing Query Performance in Distributed NoSQL Databases through Adaptive Indexing and Data Partitioning Techniques. International Journal of Creative Research Thoughts, 10(8), e812-e823.

[8] Krishna, K. (2020). Towards Autonomous AI: Unifying Reinforcement Learning, Generative Models, and Explainable AI for Next-Generation Systems. Journal of Emerging Technologies and Innovative Research, 7(4), 60-68.

[9] Murthy, P. & Mehra, A. (2021). Exploring Neuromorphic Computing for Ultra-Low Latency Transaction Processing in Edge Database Architectures. Journal of Emerging Technologies and Innovative Research, 8(1), 25-33

[10] Yalamati, S. (2024). Data privacy, compliance, and security in cloud computing for finance. In Practical Applications of Data Processing, Algorithms, and Modeling (pp. 127-144). IGI Global.

[11] Katari, A., & Ankam, M. (2022). Data Governance in Multi-Cloud Environments for Financial Services: Challenges and Solutions. Educational Research (IJMCER), 4(1), 339-353.

[12] Owoade, S. J., Uzoka, A., Akerele, J. I., & Ojukwu, P. U. (2024). Cloud-based compliance and data security solutions in financial applications using CI/CD pipelines. World Journal of Engineering and Technology Research, 8(2), 152-169.

[13] Thota, R. C. (2021). Cloud Security in Financial Services: Protecting Sensitive Data with AWS well-Architected Framework. INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT, 6(4), 1-7.

[14] Atadoga, A., Umoga, U. J., Lottu, O. A., & Sodiya, E. O. (2024). Evaluating the impact of cloud computing on accounting firms: A review of efficiency, scalability, and data security. Global Journal of Engineering and Technology Advances, 18(2), 065-074.

[15] Subramanyam, S. V. (2021). Cloud computing and business process re-engineering in financial systems: The future of digital transformation. International Journal of Information Technology and Management Information Systems (IJITMIS), 12(1), 126-143.

[16] Salako, A. O., Fabuyi, J. A., Aideyan, N. T., Selesi-Aina, O., Dapo-Oyewole, D. L., & Olaniyi, O. O. (2024). Advancing information governance in AI-driven cloud ecosystem: Strategies for enhancing data security and meeting regulatory compliance. Asian Journal of Research in Computer Science, 17(12), 66-88.

[17] Oladoyinbo, T. O., Adebiyi, O. O., Ugonnia, J. C., Olaniyi, O. O., & Okunleye, O. J. (2023). Evaluating and establishing baseline security requirements in cloud computing: an enterprise risk management approach. Asian journal of economics, business and accounting, 23(21), 222-231.

[18] Krishna, K. & Thakur, D. (2021). Automated Machine Learning (AutoML) for Real-Time Data Streams: Challenges and Innovations in Online Learning Algorithms. Journal of Emerging Technologies and Innovative Research, 8(12), f730-f739.

[19] Mehra, A. (2024). Hybrid AI Models: Integrating Symbolic Reasoning with Deep Learning for Complex Decision-Making. Journal of Emerging Technologies and Innovative Research, 11(8), f693-f704.

[20] Murthy, P. & Thakur, D. (2022). Cross-Layer Optimization Techniques for Enhancing Consistency and Performance in Distributed NoSQL Database. International Journal of Enhanced Research in Management & Computer Applications, 11(8), 35-41.

[21] Murthy, P. (2020). Optimizing Cloud Resource Allocation using Advanced AI Techniques: A Comparative Study of Reinforcement Learning and Genetic Algorithms in Multi-Cloud Environments. World Journal of Advanced Research and Reviews, 7(2), 359-369.

[22] Mehra, A. (2021). Uncertainty Quantification in Deep Neural Networks: Techniques and Applications in Autonomous Decision-Making Systems. World Journal of Advanced Research and Reviews, 11(3), 482-490.

[23] A. Gatouillat, Y. Badr, B. Massot, and E. Sejdić, "Internet of Medical Things: A Review of Recent Contributions Dealing with Cyber-Physical Systems in Medicine," IEEE Internet of Things Journal, vol. 5, no. 5, pp. 3810-3822, 2018.

[24] C. Esposito, A. De Santis, G. Tortora, H. Chang, and K. K. R. Choo, "Blockchain: A Panacea for Healthcare Cloud-Based Data Security and Privacy?," IEEE Cloud Computing, vol. 5, no. 1, pp. 31-37, 2018.

[25] M. Abdel-Basset, G. Manogaran, A. Gamal, and V. Chang, "A Novel Intelligent Medical Decision Support Model Based on Soft Computing and IoT," IEEE Internet of Things Journal, vol. 7, no. 5, pp. 4160-4170, 2020.