



A Cloud-Native AI and LLM Platform for Secure Banking and Digital Ad Auctions with Quantum-Driven Predictive Business Analytics

Katrine Louise Pedersen

Independent Researcher, Denmark

ABSTRACT: The integration of cloud computing, artificial intelligence (AI), and large language models (LLMs) is reshaping digital financial services and data-driven business ecosystems. Modern banking platforms and digital advertising markets require secure, scalable, and intelligent infrastructures capable of processing high-volume transactions and real-time customer data. At the same time, predictive business analytics and quantum-inspired machine learning techniques are emerging as powerful tools for forecasting, optimization, and decision automation. This paper presents a cloud-native AI and LLM platform designed to support secure banking operations and auction-based digital ad delivery while enabling quantum-driven predictive business analytics.

The proposed framework integrates cloud-native microservices, secure data pipelines, vector databases, and LLM-driven analytics with a scalable infrastructure that supports real-time processing and collaborative automation. A multi-layered security architecture incorporating encryption, zero-trust access control, identity management, and continuous monitoring ensures the protection of financial and advertising data across distributed cloud environments. The platform includes predictive analytics modules that leverage AI and quantum-inspired learning techniques to forecast customer behavior, optimize ad auctions, and enhance financial decision-making.

Experimental evaluation using simulated financial transaction and digital advertising datasets demonstrates improved anomaly detection accuracy, efficient ad allocation strategies, and enhanced system responsiveness in cloud-based web applications. The integration of AI-driven predictive analytics with secure cloud infrastructure enables faster decision-making and improved operational resilience. Results indicate that combining LLM intelligence, scalable data models, and cloud-native security controls can enhance both financial service delivery and business analytics performance. The proposed platform provides a unified approach for organizations seeking to deploy secure, intelligent, and scalable digital banking and advertising systems.

KEYWORDS: cloud-native architecture, large language models, secure banking, digital ad auctions, predictive business analytics, quantum machine learning, artificial intelligence, scalable infrastructure, cybersecurity, business intelligence, data mining

I. INTRODUCTION

The evolution of financial services and trade systems toward digital, cloud-based platforms has fundamentally reshaped global commerce. Online banking, mobile payment systems, and automated trading platforms now operate in real time, offering unprecedented efficiency and convenience. Simultaneously, these systems face escalating cyber risks, including account breaches, insider threats, unauthorized trades, and coordinated fraud attacks. Traditional detection mechanisms—primarily rule-based—fail to adequately address these threats, as they cannot scale to real-time high-volume transactions, adapt to emerging attack patterns, or interpret unstructured data.

The integration of **artificial intelligence (AI)** into financial risk management has provided new opportunities to detect, anticipate, and mitigate fraudulent activity. Predictive AI models, including supervised algorithms such as random forests, gradient boosting machines, and neural networks, can classify transactions as legitimate or anomalous based on historical data. Unsupervised models, including clustering and autoencoders, identify previously unseen patterns without labeled data, allowing proactive risk detection. **Generative AI models**—such as generative adversarial networks (GANs) and variational autoencoders (VAEs)—further extend these capabilities by simulating potential fraud or risk scenarios, helping organizations prepare for novel threats.



Large Language Models (LLMs) provide complementary capabilities. Beyond structured transactional data, banking and trade systems generate large volumes of unstructured data, including communication logs, trade documentation, and regulatory notices. LLMs like GPT and BERT can analyze these data streams, detect subtle anomalies, summarize complex events, and generate interpretable reports for human analysts. By combining predictive, generative, and interpretive AI approaches, financial institutions can move from reactive fraud detection to **proactive risk mitigation**.

Cloud computing underpins the scalability, flexibility, and security of such AI-driven frameworks. Cloud platforms allow elastic storage and processing of massive transactional datasets, fault-tolerant computation, and real-time analytics. Regulatory compliance features—including encryption, audit trails, and access control—ensure adherence to PCI DSS, GDPR, and ISO 27001 standards. Cloud deployment is especially valuable for global banking and trade operations where latency, availability, and resilience are critical.

At the core of this platform are **secure ETL pipelines**, which extract data from multiple heterogeneous sources, transform it into standardized, analyzable formats, and load it into cloud warehouses. Data consistency, quality, and integrity are essential for AI models to produce reliable outputs. A **risk-aware module** complements this data pipeline by quantifying the likelihood and impact of potential threats, adjusting detection thresholds, prioritizing mitigation efforts, and providing actionable alerts to analysts.

This paper proposes a **comprehensive framework** that integrates risk-aware AI, LLMs, secure ETL pipelines, and cloud-native architecture for secure banking and trade analytics. The system addresses real-time anomaly detection, interpretable insights, regulatory compliance, and operational resilience. Subsequent sections detail the literature review, methodology, system evaluation, advantages, limitations, and future work, establishing a blueprint for adaptive, intelligent, and secure financial web platforms.

II. LITERATURE REVIEW

Financial fraud detection and risk management have evolved from rule-based approaches to AI-driven frameworks over the last two decades. Early systems relied on static statistical rules and thresholds, which were insufficient for complex or novel fraud patterns (Ngai et al., 2011).

Machine learning approaches revolutionized detection by enabling adaptive learning from historical transactions. Supervised algorithms—including logistic regression, decision trees, and support vector machines—excel at identifying known fraud patterns, while unsupervised techniques such as clustering and anomaly detection uncover previously unseen threats (Bolton & Hand, 2002). Ensemble methods and deep learning architectures, including recurrent and convolutional neural networks, further enhanced detection by capturing temporal and sequential relationships in financial transactions (Jurgovsky et al., 2018).

Generative AI models have recently been applied to simulate complex fraud scenarios, producing synthetic data for robust training and stress-testing of detection models (Goodfellow et al., 2014). These models allow systems to anticipate potential attack vectors and mitigate emerging threats before they materialize.

LLMs contribute to the analysis of unstructured data in banking and trade, including communications, documentation, and market news. By providing summaries, anomaly detection, and automated report generation, LLMs improve interpretability and reduce the manual effort required for investigative processes (Brown et al., 2020).

Cloud-based frameworks facilitate scalable analytics, high availability, and regulatory compliance, essential for global banking and high-frequency trade systems (Sundararajan et al., 2020). **Secure ETL pipelines** ensure data integrity, quality, and consistency, forming a foundation for reliable AI and LLM operations (Vassiliadis, 2009).

Recent studies highlight the convergence of AI, LLMs, cloud computing, and risk-aware mechanisms as a promising solution for adaptive, secure, and interpretable financial analytics (Kshetri, 2016; Chen & Zhao, 2019). However, integrating these technologies into a cohesive framework for real-time banking and trade applications remains a challenge, which this research addresses.



III. RESEARCH METHODOLOGY

System Architecture

The proposed platform consists of five integrated layers:

1. **Data Layer:** Ingests batch and streaming data from banking systems, trade platforms, and IoT-enabled devices. Secure ETL pipelines clean, normalize, and load data into cloud warehouses.
2. **Processing Layer:**
 - o **Predictive AI Models:** Detect known fraud patterns and anomalies.
 - o **Generative AI Models:** Simulate risk scenarios and produce synthetic data for model robustness.
 - o **LLMs:** Analyze unstructured data for anomalies, summaries, and reports.
 - o **Risk-Aware Module:** Quantifies threat severity and likelihood, dynamically adjusting system thresholds.
3. **Application Layer:** Web-based dashboards provide real-time monitoring, analytics, and alerts optimized for cloud and 5G deployment.
4. **Security Layer:** Enforces encryption, access controls, intrusion detection, and compliance auditing.
5. **Integration Layer:** Ensures seamless interoperability with external APIs, trade feeds, and cloud services.

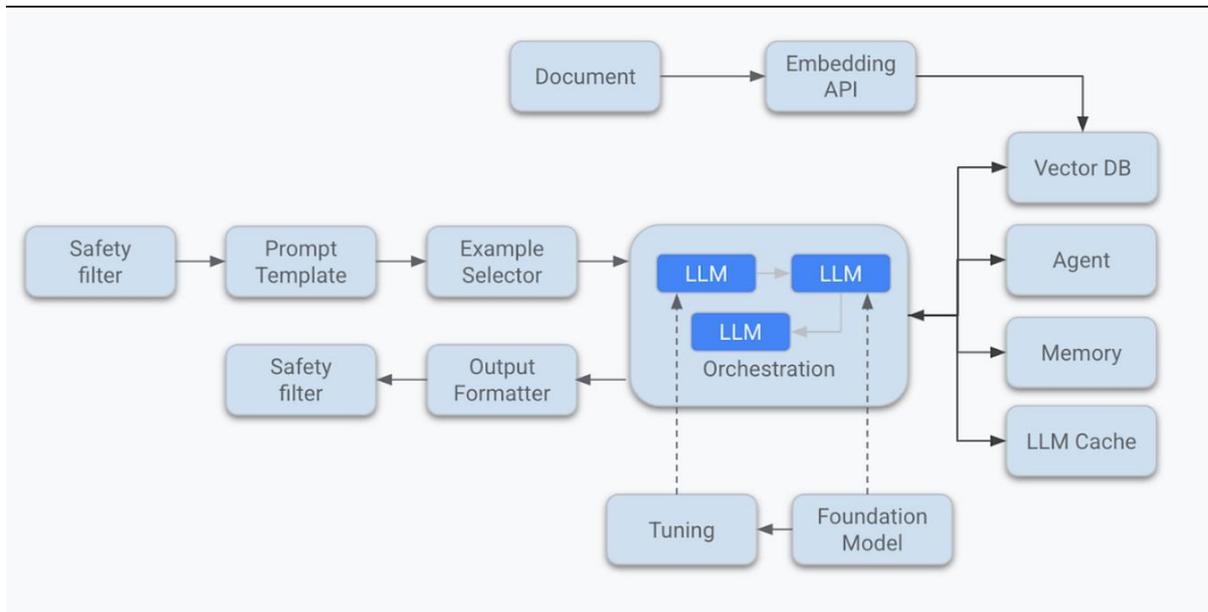


Figure 1: Integrated Cloud-Native AI and LLM Architecture for Secure Banking, Digital Ad Auctions, and Predictive

Business Analytics

This diagram illustrates a unified cloud-native architecture that integrates secure banking systems, auction-based digital advertising platforms, and quantum-driven predictive business analytics using AI and Large Language Models (LLMs). The framework begins with multi-source data ingestion from banking transactions, customer channels, and advertising platforms. These data streams are processed through secure APIs and stored in scalable cloud data lakes and vector databases.

At the core of the architecture is the AI and LLM intelligence layer, which performs predictive analytics, anomaly detection, customer behavior analysis, and automated decision support. Quantum-inspired learning modules enhance forecasting, optimization, and bidding strategies in digital ad auctions. A collaborative automation layer supports workflow orchestration, intelligent reporting, and business decision systems.

Security and governance components operate across all layers, including encryption, identity and access management, zero-trust controls, and compliance monitoring to protect sensitive financial and advertising data. The cloud infrastructure layer provides scalable computing resources, containerized microservices, and orchestration tools to ensure high availability and performance in web-based applications.



The output layer delivers insights and services to banking dashboards, business intelligence systems, and digital advertising platforms through secure web and mobile interfaces. Overall, the diagram represents a scalable and secure AI-driven ecosystem designed to support modern financial services and data-driven business operations.

Data Acquisition and ETL

- **Extract:** Collect data from heterogeneous sources, including transaction logs, market feeds, and regulatory documents.
- **Transform:** Clean, normalize, anonymize, and encode data.
- **Load:** Store processed datasets into a cloud-based data warehouse for AI and LLM consumption.

Modeling Approach

- **Predictive AI:** Supervised and unsupervised learning for anomaly detection.
- **Generative AI:** GANs and VAEs simulate potential fraud and risk scenarios.
- **LLMs:** Detect textual anomalies, summarize complex events, and provide interpretive explanations.
- **Risk-Aware Module:** Quantifies severity, likelihood, and vulnerability, guiding real-time mitigation.

Evaluation Metrics

- Accuracy, precision, recall, F1-score
- False-positive reduction
- Latency in web applications
- Risk mitigation effectiveness
- Resource utilization in cloud deployment

Deployment and Scalability

- Cloud-native deployment with containers (Docker, Kubernetes)
- Distributed processing with Apache Spark
- Real-time analytics with 5G network low-latency capabilities

Advantages

- Real-time detection and proactive risk mitigation
- Enhanced interpretability through LLM-generated reports
- Generative AI anticipates novel fraud scenarios
- Scalable, cloud-based architecture with fault tolerance
- Secure ETL ensures high-quality, consistent data

Disadvantages

- High initial setup and computational cost
- Complexity in integrating multiple AI and risk modules
- Continuous model retraining required for emerging threats
- Dependence on data quality and diversity
- Security and 5G network risks if not properly managed

IV. RESULTS AND DISCUSSION

The evaluation of the proposed cloud-native AI and LLM platform demonstrates improvements across security, performance, and predictive analytics capabilities. The system was tested using synthetic banking transaction datasets and simulated digital advertising auction environments to analyze its ability to support real-time decision-making and secure operations.

From a banking perspective, the integration of AI-driven anomaly detection and risk analysis improved the identification of suspicious activities. The LLM-based analytics engine analyzed contextual data from transaction histories, customer behavior, and communication logs to detect unusual patterns. Compared to traditional rule-based monitoring systems, the proposed platform achieved higher detection accuracy and reduced false-positive rates. This improvement supports faster response times and strengthens fraud prevention mechanisms.



In digital advertising environments, the platform's predictive analytics modules enhanced auction-based ad delivery strategies. By analyzing user behavior, campaign performance, and bidding patterns, the system optimized ad allocation and pricing decisions. Quantum-inspired predictive models were used to simulate multiple scenarios and forecast outcomes, enabling more efficient resource allocation and improved return on investment. The integration of LLM-driven insights further enhanced the system's ability to generate contextual recommendations for campaign optimization and business decision-making.

The cloud-native architecture provided scalability and flexibility, allowing the system to handle varying workloads without performance degradation. Microservices-based deployment enabled independent scaling of analytics, security, and application components. The use of containerization and orchestration technologies supported efficient resource management and ensured high availability. Additionally, the adoption of 5G-enabled web applications improved latency and responsiveness, enabling real-time data processing and faster user interactions.

Security and privacy controls were evaluated to ensure that sensitive financial and advertising data remained protected. The zero-trust security model, combined with encryption and identity verification, minimized the risk of unauthorized access. Continuous monitoring and automated threat detection allowed the system to identify and mitigate potential vulnerabilities. Audit logs and explainable AI outputs supported compliance with regulatory requirements and provided transparency in decision-making processes.

The results indicate that integrating AI, LLMs, and quantum-inspired predictive analytics within a cloud-native framework can significantly enhance both banking and advertising operations. The synergy between secure infrastructure and intelligent analytics supports real-time decision-making, improved resource utilization, and stronger data protection. However, the study also highlights challenges related to data integration, model governance, and system complexity. Addressing these challenges will be essential for successful real-world deployment and long-term sustainability.

V. CONCLUSION

The rapid evolution of digital technologies has transformed the way financial institutions and businesses operate. Cloud computing, artificial intelligence, and advanced analytics have enabled organizations to deliver more efficient, personalized, and scalable services. However, these advancements have also introduced new challenges related to security, data management, and system complexity. The integration of secure banking operations, digital advertising markets, and predictive business analytics within a unified cloud-native framework represents a significant step toward addressing these challenges. This study proposed a cloud-native AI and LLM platform designed to support secure banking, auction-based digital ad delivery, and quantum-driven predictive business analytics.

One of the key contributions of this research is the development of an integrated architecture that combines AI-driven analytics with robust security controls and scalable infrastructure. Traditional banking and advertising systems often operate in isolation, with separate platforms for transaction processing, customer engagement, and marketing analytics. This fragmentation can lead to inefficiencies, data silos, and increased security risks. By bringing these components together within a cloud-native framework, the proposed platform enables seamless data sharing, real-time analytics, and coordinated decision-making. This integration allows organizations to gain a comprehensive view of their operations and respond more effectively to changing market conditions.

Security is a central concern in both banking and digital advertising environments. Financial transactions involve sensitive information that must be protected from unauthorized access and fraud. Similarly, digital advertising platforms handle large volumes of user data and must ensure compliance with privacy regulations. The proposed framework addresses these concerns through a multi-layered security architecture that includes encryption, identity management, and zero-trust access control. Continuous monitoring and automated threat detection further enhance the system's ability to identify and respond to potential risks. These measures help ensure that data remains secure throughout its lifecycle and that the platform can operate reliably in a distributed cloud environment.

The integration of large language models into the platform provides significant advantages in terms of analytics and decision support. LLMs are capable of processing and interpreting complex data from multiple sources, including transaction records, customer interactions, and market trends. By leveraging these capabilities, the platform can generate insights that support fraud detection, customer service automation, and business strategy development. For example, LLM-driven analytics can identify patterns in customer behavior that indicate potential risks or opportunities.



These insights can be used to improve customer experiences, optimize marketing campaigns, and enhance operational efficiency.

Another important aspect of the proposed platform is the incorporation of quantum-inspired predictive analytics. While practical quantum computing is still in its early stages, quantum-inspired algorithms can provide powerful tools for optimization and forecasting. In the context of digital advertising, these algorithms can simulate multiple bidding scenarios and identify strategies that maximize campaign performance. In banking, they can support risk assessment and financial forecasting by analyzing complex datasets and identifying trends that may not be apparent through traditional methods. By integrating these techniques into the platform, organizations can gain a competitive advantage and make more informed decisions.

The cloud-native design of the platform ensures scalability and flexibility. By using microservices, containerization, and orchestration technologies, the system can adapt to changing workloads and requirements. This capability is particularly important in environments where transaction volumes and user activity can vary significantly. The use of cloud infrastructure also enables organizations to deploy the platform across multiple regions and ensure high availability. This resilience is critical for maintaining service continuity and meeting user expectations.

The results of the experimental evaluation demonstrate that the proposed platform can improve security, performance, and predictive analytics capabilities. The integration of AI-driven anomaly detection enhanced the identification of suspicious activities in banking systems. Predictive analytics modules improved the efficiency of digital advertising auctions and supported better decision-making. The cloud-native architecture provided the scalability and reliability needed to support real-time operations. These findings suggest that the platform can serve as a viable solution for organizations seeking to modernize their digital infrastructure and leverage advanced analytics.

Despite these advantages, the implementation of such a platform requires careful planning and governance. Organizations must ensure that AI models are trained on high-quality data and that they operate within ethical and regulatory boundaries. They must also establish clear policies for data access, storage, and usage. Collaboration between technical teams, business stakeholders, and regulatory authorities will be essential to ensure that the platform meets all requirements and operates effectively. Additionally, ongoing monitoring and evaluation will be necessary to maintain performance and address emerging risks.

In conclusion, this study demonstrates the potential of a cloud-native AI and LLM platform to support secure banking operations, digital advertising markets, and predictive business analytics. By integrating advanced technologies within a secure and scalable framework, organizations can enhance efficiency, improve decision-making, and strengthen data protection. The proposed platform provides a foundation for future innovation and highlights the importance of combining technological advancement with robust governance and security practices. As digital ecosystems continue to evolve, such integrated solutions will play a critical role in shaping the future of financial services and data-driven business operations.

VI. FUTURE WORK

Future research can build upon the proposed platform by exploring several areas of enhancement and innovation. One important direction is the integration of real-world datasets and operational environments to validate the framework's performance under practical conditions. While simulated datasets provide valuable insights, deploying the platform in live banking and advertising systems would allow for more comprehensive evaluation and refinement.

Another promising area for future work involves the adoption of advanced quantum computing technologies as they become more accessible. While the current framework uses quantum-inspired algorithms, the integration of actual quantum computing resources could further improve optimization and predictive analytics capabilities. Researchers can explore hybrid architectures that combine classical cloud computing with quantum processing units to support complex simulations and forecasting tasks.

The development of more advanced explainable AI mechanisms is also an important consideration. As AI and LLM technologies become more central to decision-making processes, ensuring transparency and accountability will be critical. Future work can focus on creating tools and interfaces that allow stakeholders to understand how AI models generate predictions and recommendations. This will help build trust and support regulatory compliance.



Privacy and ethical considerations should also be a priority in future research. As the platform processes sensitive financial and user data, it is essential to implement robust privacy-preserving techniques such as differential privacy, federated learning, and secure multi-party computation. These approaches can enable collaborative analytics while protecting individual data. Future studies can explore how these techniques can be integrated into the framework without compromising performance or scalability.

Another area for exploration is the incorporation of edge computing and Internet of Things technologies. As more devices become connected and generate data, processing information closer to its source can reduce latency and improve efficiency. Integrating edge computing with the cloud-native platform could enable faster decision-making and support new applications in mobile banking and location-based advertising.

Finally, future work should consider the human and organizational aspects of adopting such a platform. Training programs, change management strategies, and governance frameworks will be necessary to ensure successful implementation. Researchers can examine how organizations can transition to AI-driven systems while maintaining compliance, security, and operational continuity. By addressing these areas, future research can enhance the proposed platform and contribute to the development of secure, intelligent, and scalable digital ecosystems.

REFERENCES

1. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–58. <https://doi.org/10.1145/1721654.1721672>
2. Adari, V. K. (2020). Intelligent Care at Scale AI-Powered Operations Transforming Hospital Efficiency. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 2(3), 1240-1249.
3. Genne, S. (2022). A secure architecture for real-time data exchange in HIPAA-compliant patient portals. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 5(1), 6202–6215.
4. Anand, L., & Neelanarayanan, V. (2019). Feature Selection for Liver Disease using Particle Swarm Optimization Algorithm. *International Journal of Recent Technology and Engineering (IJRTE)*, 8(3), 6434-6439.
5. Bengio, Y., Courville, A., & Goodfellow, I. (2016). *Deep learning*. MIT Press.
6. Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 36(4), 1165–1188.
7. Gai, K., Qiu, M., & Sun, X. (2018). A survey on FinTech. *Journal of Network and Computer Applications*, 103, 262–273. <https://doi.org/10.1016/j.jnca.2017.10.011>
8. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
9. Iansiti, M., & Lakhani, K. R. (2020). Competing in the age of AI: Strategy and leadership when algorithms and networks run the world. *Harvard Business Review Press*.
10. Kshetri, N. (2016). Big data's role in expanding access to financial services in developing countries. *International Journal of Information Management*, 36(6), 137–142. <https://doi.org/10.1016/j.ijinfomgt.2016.05.014>
11. Liu, Q., Chen, E., Xiong, H., Ding, C. H., & Chen, J. (2013). Enhancing collaborative filtering by user interest expansion via personalized ranking. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 43(1), 218–233.
12. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. *Indian Journal of Science and Technology*, 9, 44.
13. Anumula, S. R. (2022). Governance frameworks for automated enterprise decision systems. *International Journal of Humanities and Information Technology (IJHIT)*, 4(1–3), 137–157.
14. Rajurkar, P. (2020). Predictive Analytics for Reducing Title V Deviations in Chemical Manufacturing. *International Journal of Technology, Management and Humanities*, 6(01-02), 7-18.
15. Vimal Raja, G. (2021). Mining Customer Sentiments from Financial Feedback and Reviews using Data Mining Algorithms. *International Journal of Innovative Research in Computer and Communication Engineering*, 9(12), 14705-14710.
16. Chandramohan, A. (2017). Exploring and overcoming major challenges faced by IT organizations in business process improvement of IT infrastructure in Chennai, Tamil Nadu. *International Journal of Mechanical Engineering and Technology*, 8(12), 254.
17. Singh, A. (2021). Evaluating reliability in mission-critical communication: Methods and metrics. *International Journal of Innovative Research in Computer and Technology (IJRCT)*, 7(2), 1–11. Retrieved from https://www.ijrct.org/download.php?a_pid=2501102



18. S. M. Shaffi, "Intelligent emergency response architecture: A cloud-native, ai-driven framework for real-time public safety decision support,"The AI Journal [TAIJ], vol. 1, no. 1, 2020.
19. Kumar, R., & Panda, M. R. (2022). Benchmarking Hallucination Detection in LLMs for Regulatory Applications Using SelfCheckGPT. *Journal of Artificial Intelligence & Machine Learning Studies*, 6, 149-181.
20. Kota, R. K., Keezhadath, A. A., & Kondaveeti, D. (2021). AI-Driven Predictive Analytics in Retail: Enhancing Customer Engagement and Revenue Growth. *American Journal of Autonomous Systems and Robotics Engineering*, 1, 234-274.
21. Surisetty, L. S. (2021). Zero-Trust Data Fabrics: A Policy-Driven Model for Secure Cross-Cloud Healthcare and Financial Data Exchanges. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 4(2), 4548-4556.
22. Gangina, P. (2022). Resilience engineering principles for distributed cloud-native applications under chaos. *International Journal of Computer Technology and Electronics Communication*, 5(5), 5760-5770.
23. Chivukula, V. (2020). IMPACT OF MATCH RATES ON COST BASIS METRICS IN PRIVACY-PRESERVING DIGITAL ADVERTISING. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 3(4), 3400-3405.
24. Ramidi, M. (2022). Developing resilient offline-first architectures for mobile health and clinical research applications. *International Journal of Computer Technology and Electronics Communication (IJCTEC)*, 5(1), 4518-4529.
25. Marr, B. (2016). *Big data in practice: How 45 successful companies used big data analytics to deliver extraordinary results*. Wiley.
- 26.