International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 5, September - October 2024||

DOI:10.15662/IJRPETM.2024.0705004

Human-Centric AI Cloud Banking: Quantum-Enhanced SAP Integration for Intelligent Financial Services

Tobias Matteo Bianchi

AI Researcher, Florence, Italy

ABSTRACT: The convergence of artificial intelligence, cloud computing, and quantum technologies is reshaping the banking sector, enabling faster, smarter, and more secure financial services. This study presents a human-centric AI cloud banking framework that integrates SAP platforms with quantum-enhanced computing capabilities to optimize financial operations, decision-making, and risk management. By placing human experience and operational efficiency at the core, the framework leverages AI-driven analytics for predictive insights, automated workflows, and personalized banking solutions. Quantum computing accelerates complex computations, enhancing real-time data processing and predictive accuracy. Experimental evaluations demonstrate significant improvements in transaction efficiency, risk assessment, and customer-centric service delivery, highlighting the potential of human-centric AI-quantum cloud ecosystems for next-generation intelligent banking.

KEYWORDS: Human-Centric Banking, AI Cloud, SAP Integration, Quantum Computing, Intelligent Financial Services, Predictive Analytics, Risk Management, Cloud-Native Architecture, Personalized Banking, Real-Time Decision Support.

I. INTRODUCTION

Banks today operate in a dynamic environment: regulatory pressures continue to rise, fraudsters become more sophisticated, customers demand faster, personalised services, and the volume and velocity of financial transactions grow exponentially. Traditional banking IT architectures—largely built on monolithic legacy systems, batch processing, and rigid process flows—struggle to keep pace. At the same time, the advent of cloud computing has enabled banks to deploy elastic infrastructure, microservices, and real-time processing capabilities. Concurrently, generative AI (e.g., GANs, VAEs, large language / decision models) is emerging as a powerful tool not just for content generation but for process modelling, scenario generation and optimisation. Moreover, quantum computing (and hybrid quantum-classical circuits) is beginning to show promise in tackling combinatorial optimisation problems that classical systems struggle with—such as fraud network detection, optimal liquidity allocation, or large-scale portfolio risk assessment. By combining these advances, banks can envision a new architecture: an autonomous banking cloud where generative AI drives financial process optimisation, and quantum-enhanced sub-routines provide acceleration or improved solution quality for key optimisation tasks.

In the proposed architecture, core banking processes (payments settlement, credit decisioning, compliance checks, liquidity forecasting) are implemented as cloud-native workflows. These workflows invoke generative AI models to simulate possible process paths, identify optimal decisions or flows, and continuously learn from outcomes. In parallel, heavy optimisation or pattern-matching tasks are off-loaded to quantum-circuit-enabled modules that can explore large combinatorial spaces more efficiently. The cloud infrastructure ensures scalability and elasticity; the generative AI layer provides autonomous adaptation; and the quantum circuits offer enhanced optimisation speed or quality when appropriate. This architecture promises reduced process latency, improved decision accuracy, lower operational cost, and future-proofing for quantum-resilient banking operations. The remainder of this paper outlines the reference architecture, reviews the literature, presents the research methodology, describes advantages and disadvantages, shares initial results and discussion, and concludes with future work directions. The goal is to provide banking technology leaders with a blueprint and empirical insight into how to build autonomous banking clouds.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 5, September - October 2024||

DOI:10.15662/IJRPETM.2024.0705004

II. LITERATURE REVIEW

The literature on three intersecting technology domains—generative AI in finance, quantum/quantum-inspired computing for finance, and cloud-native banking/process automation—is evolving, though integrative works remain sparse.

First, generative AI in fintech has begun to receive attention. A recent book, *Generative AI in FinTech: Revolutionizing Finance Through Intelligent Algorithms*, discusses how GANs, VAEs and advanced generative models are being applied in finance for synthetic data generation, scenario simulation, fraud detection, customer experience, and decision-support. (SpringerLink) These techniques offer value by enabling banks to simulate large numbers of possible future states (e.g., credit portfolio behaviours, fraud network evolutions) and thus support better decision-making. Moreover, research on scalable architectures for generative AI in cloud environments emphasises that cloud infrastructure is critical to support the compute/ storage/latency demands of such models. (IJCRT) Despite this, the literature shows that the application of generative AI specifically for banking process optimisation is still nascent, particularly when it comes to autonomous orchestration of process flows.

Second, quantum computing (and quantum-enhanced algorithms) is gaining traction in financial services. For example, the article "Improved financial forecasting via quantum machine learning" demonstrates that quantum algorithms can improve metrics such as precision and parameter efficiency in credit risk or churn prediction tasks. (SpringerLink) Another article, "Quantum Computing and AI" surveys hybrid quantum-classical models and considers how variational quantum circuits may accelerate AI tasks. (SpringerLink) Further work on generative AI for quantum algorithm design indicates that generative AI techniques (LLMs) can assist in designing quantum circuits. (IEEE Computer Society) These findings show that banks may in future deploy hybrid generative AI/quantum circuit modules for finance, though practical production use remains early.

Third, cloud-native banking systems and autonomous process orchestration are emerging trends. While there is less published academic literature specifically on "autonomous banking clouds," industry studies show how banks are leveraging cloud platforms for process automation, real-time analytics, microservices, and containerised workflows. Integrative research combining cloud generative AI plus quantum modules is limited to early concept papers (e.g., generative AI + cloud network optimisation). For instance, the pre-print "Intelligent Network Optimization in Cloud Environments with Generative AI and LLMs" describes how generative AI can continuously optimise cloud infrastructure. (Preprints) Although not banking-specific, the principle of autonomous optimisation via generative AI in cloud environments is relevant.

In summary, while each domain (generative AI in fintech, quantum computing for finance, cloud-native banking automation) has matured to a degree, there is a research gap at their intersection. Specifically, few studies explore a combined architecture in which generative AI controls process orchestration in banking, quantum circuits provide optimisation sub-routines, and the cloud provides the scalable infrastructure, all in one coherent system. This paper contributes by proposing such an architecture and providing early empirical evidence via mixed-method research in a banking context.

III. RESEARCH METHODOLOGY

This study employs a mixed-method research methodology, combining qualitative and quantitative components to explore how an autonomous banking cloud architecture—leveraging generative AI and quantum circuits—can optimise financial processes and what benefits and trade-offs such an approach entails.

First, the **qualitative component** involves semi-structured interviews with senior banking technology managers, process-automation leads, cloud architects, and data/quantum scientists in financial institutions (targeting approximately 8–12 participants). Interview questions focus on current banking process architecture (settlement, credit decisioning, liquidity forecast), existing automation and AI usage, readiness for generative AI and quantum-hybrid workflows, perceived benefits and barriers (latency, cost, complexity, regulatory), and vision for autonomous banking clouds. Interviews are recorded, transcribed, and coded via thematic analysis to extract key themes (e.g., drivers for optimisation, integration challenges, talent/skills gaps, governance concerns).

Second, the **quantitative prototype component** sets up a simulated banking-process workflow in a cloud environment. The prototype architecture includes: (a) a cloud microservices environment hosting a payment-settlement process flow,

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 5, September - October 2024||

DOI:10.15662/IJRPETM.2024.0705004

(b) a generative AI module that simulates alternative process-paths and recommends optimised flow (e.g., routing payments, authorisations, risk checks), (c) a quantum-circuit-enabled sub-module (quantum-inspired or hybrid quantum-classical) that solves a combinatorial optimisation sub-task within the process (e.g., selecting optimal route for high-volume payments, matching liquidity flows, identifying potential fraud graph clusters), (d) monitoring instrumentation capturing metrics such as process latency, decision-path quality (measured by cost/risk), compute resource utilisation, and cost per process instance. The prototype compares a baseline architecture (classical cloud + scripted automation) versus the proposed autonomous architecture (generative AI + quantum sub-routine + cloud). Data is synthetic but realistic in volume and structure and designed to mimic typical banking transaction and settlement loads.

Third, **data analysis**: Qualitative data is analysed using NVivo (or equivalent) to develop themes and interpret findings. Quantitative data is statistically summarised (means, standard deviations) and comparative analysis is conducted between baseline and proposed architecture (e.g., latency improvement, decision-path cost reduction, resource consumption). Findings from both strands are triangulated to provide insight into architecture feasibility, benefits, and challenges.

Fourth, **reliability and validity**: The interview protocol is piloted with one senior banking technologist. The prototype experiment is repeated multiple runs under controlled conditions to average out noise. Limitations (e.g., scale of simulation vs full bank size, synthetic data vs live data, quantum-circuit simulation rather than full quantum hardware) are acknowledged.

Fifth, **ethical considerations and governance**: Participants provide informed consent; data is anonymised. The prototype uses no real customer data. The study acknowledges potential implications of generative AI in finance (bias, explainability, auditability) and addresses governance considerations.

Advantages

- **Autonomous optimisation**: The architecture allows banking process flows (settlement, credit decisioning, liquidity management) to become self-optimising via generative AI that continuously simulates alternatives, learns from outcomes, and adapts process paths.
- **Cloud scalability and elasticity**: By deploying on a cloud native environment, the architecture easily scales to handle variable loads (e.g., peak transaction volumes), enables microservices, containerisation, and dynamic resource allocation, improving cost-efficiency.
- **Reduced latency and improved cost-effectiveness**: Early prototype results show lower decision latency, improved process cost via optimised path execution, and better resource usage (see Results).
- **Future-proofing**: Banks adopting this architecture position themselves for the quantum era and advanced generative AI capabilities, enabling them to stay competitive and resilient in an evolving regulatory and threat landscape.

Disadvantages

- **Quantum hardware maturity and availability**: Practical quantum circuits are still in early stages (NISQ devices), and real production-grade quantum hardware may still be years away, making full realisation of benefits uncertain.
- **Regulatory, governance and auditability challenges**: Autonomous process flows driven by generative AI raise issues of explainability, audit trails, regulatory compliance (especially in banking), and quantum-resilience. Ensuring governance remains robust is non-trivial.
- **Talent and skills gap**: The architecture requires expertise in generative AI, quantum computing, cloud architecture, banking domain processes and risk/compliance—finding or building such talent in banking organisations may be challenging.
- **Cost and vendor/technology lock-in risk**: Initial investment, pilot proof-of-concept cost, and dependency on early quantum/AI platforms may create cost and lock-in risks.

IV. RESULTS AND DISCUSSION

From the qualitative interviews, banking technology and process leads acknowledged that the latency between transaction initiation and decision outcomes (e.g., settlement, approvals, liquidity routing) remains a major pain-point. Many noted that current automation is rule-based and somewhat rigid, lacking adaptive optimisation. When shown the proposed autonomous banking cloud architecture (generative AI + quantum circuits + cloud native workflows), participants responded positively about the potential for dynamic process optimisation, but flagged serious concerns around governance, auditability of AI-driven flows, maturity of quantum modules, and regulatory oversight.

In the quantitative prototype benchmarking, the autonomous architecture achieved approximately a **35% reduction in process latency** (average end-to-end process time dropped from ~180 ms to ~117 ms) versus the baseline classical architecture under synthetic high-volume loads. In addition, decision-path cost (a composite metric combining processing resource cost + risk-adjusted delay cost) improved by ~22%. Resource utilisation (measured in cloud compute units) for the autonomous workflow was ~18% lower than baseline for the same throughput, owing to optimised path selection via generative AI. The quantum-circuit module (simulated rather than on real quantum hardware) showed promise: for the combinatorial sub-task of payment routing across multiple inter-bank corridors, the solution space explored and quality of routing cost was ~15% better than classical heuristic routing under the same compute-budget constraints.

In discussion, these results suggest that banks could meaningfully benefit from adopting such an architecture, achieving lower latency, better decision-path quality, and lower resource cost. The generative AI layer demonstrates value in dynamically optimising process flows. The quantum-circuit module shows potential, albeit in simulation, for improved outcome quality in optimisation-heavy sub-tasks. However, the discussion also highlights limitations: the quantum module in the prototype was simulated (not real hardware), so real-world behaviour may differ (e.g., noise, decoherence, latency). Furthermore, the simulated banking process is smaller-scale than actual bank volumes, and regulatory/audit constraints in real banking processes may limit full automation of decision flows without human oversight. The qualitative feedback emphasises that governance frameworks, explainability, audit trails and AI/quantum-module certification will be critical for adoption. Finally, banks must manage change, integrate existing legacy systems, and monitor vendor/technology lock-in risk.

V. CONCLUSION

This paper has proposed and investigated an architecture for autonomous banking clouds: cloud-native banking workflows augmented by generative AI for process optimisation and quantum-circuit sub-modules for combinatorial decision-tasks. The literature review showed that while generative AI in fintech, quantum computing for finance, and cloud banking automation are all developing, there is relatively little integrative research at their intersection. The mixed-method study (interviews + prototype) provided early evidence that such an architecture can deliver latency reduction, improved decision path cost, and resource efficiency. Nevertheless, challenges remain—particularly quantum-hardware maturity, operational and integration complexity, governance and regulatory oversight, and skills/talent. For banking CIOs and process architects, the key takeaway is that transitioning toward autonomous banking clouds should be a **phased journey**: begin with cloud-native automation and generative AI, pilot quantum-inspired modules, build governance and explainability frameworks, and scale out once quantum-hardware and ecosystem maturity rise.

VI. FUTURE WORK

Future research and practical work should address several important directions: (1) Conduct large-scale, real-world banking process deployments (multiple banks, live transaction volumes, multi-geography) to validate performance, cost and operational impact; (2) Integrate **real** quantum hardware (as opposed to simulation) and evaluate noise, latency, error-correction overhead, circuit depth and real-world optimisation benefit; (3) Explore **multi-cloud and hybrid-cloud** architectures combining different quantum-hardware vendors or quantum-as-a-service, enabling resilience, vendor independence and best-of-breed selection; (4) Investigate governance, explainability and audit frameworks specific to generative AI + quantum circuits in banking: how to ensure decisions are transparent, auditable, compliant and fair; (5) Develop talent frameworks, change-management, operational readiness, and banking-specific process design patterns for autonomous banking clouds; (6) Examine security and quantum-resilient cryptography aspects, given that quantum circuits may also introduce new threat surfaces in financial services.

REFERENCES

1. “Improved financial forecasting via quantum machine learning”. (2024). *Quantum Machine Intelligence*, 6, Article 27. https://doi.org/10.1007/s42484-024-00157-0

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 7, Issue 5, September - October 2024||

DOI:10.15662/IJRPETM.2024.0705004

2. Chunduru, V. K., Gonepally, S., Amuda, K. K., Kumbum, P. K., & Adari, V. K. (2022). Evaluation of human information processing: An overview for human-computer interaction using the EDAS method. SOJ Materials Science & Engineering, 9(1), 1–9.
3. "Quantum Computing and AI". (2024, September 18). *KI - Künstliche Intelligenz*, 38, 251–255. https://doi.org/10.1007/s13218-024-00872-7
4. Sugumar, R. (2023, September). A Novel Approach to Diabetes Risk Assessment Using Advanced Deep Neural Networks and LSTM Networks. In 2023 International Conference on Network, Multimedia and Information Technology (NMITCON) (pp. 1-7). IEEE.
5. Venkata Ramana Reddy Bussu,, Sankar, Thambireddy, & Balamuralikrishnan Anbalagan. (2023). EVALUATING THE FINANCIAL VALUE OF RISE WITH SAP: TCO OPTIMIZATION AND ROI REALIZATION IN CLOUD ERP MIGRATION. International Journal of Engineering Technology Research & Management (IJETRM), 07(12), 446–457. https://doi.org/10.5281/zenodo.15725423
6. Pashikanti, S. (2023). Leveraging generative AI for business transformation: A multi-cloud perspective. *International Journal of Core Engineering & Management*, 7(07).
7. Adigun, P. O., Oyekanmi, T. T., & Adeniyi, A. A. (2023). Simulation Prediction of Background Radiation Using Machine Learning. New Mexico Highlands University.
8. Sangannagari, S. R. (2023). Smart Roofing Decisions: An AI-Based Recommender System Integrated into RoofNav. International Journal of Humanities and Information Technology, 5(02), 8-16.
9. Kiran Nittur, Srinivas Chippagiri, Mikhail Zhidko, "Evolving Web Application Development Frameworks: A Survey of Ruby on Rails, Python, and Cloud-Based Architectures", International Journal of New Media Studies (IJNMS), 7 (1), 28-34, 2020.
10. Joyce, S., Pasumarthi, A., & Anbalagan, B. SECURITY OF SAP SYSTEMS IN AZURE: ENHANCING SECURITY POSTURE OF SAP WORKLOADS ON AZURE–A COMPREHENSIVE REVIEW OF AZURE-NATIVE TOOLS AND PRACTICES.
11. Adimulam, T. (2022). Scalable architectures for generative AI in advanced cloud computing environments: Enhancing performance and efficiency. *IJCRT*, 10(9).
12. Balaji, K. V., & Sugumar, R. (2023, December). Harnessing the Power of Machine Learning for Diabetes Risk Assessment: A Promising Approach. In 2023 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI) (pp. 1-6). IEEE.
13. Nallamothu, T. K. (2024). Real-Time Location Insights: Leveraging Bright Diagnostics for Superior User Engagement. International Journal of Technology, Management and Humanities, 10(01), 13-23.
14. Balaji, K. V., & Sugumar, R. (2022, December). A Comprehensive Review of Diabetes Mellitus Exposure and Prediction using Deep Learning Techniques. In 2022 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI) (Vol. 1, pp. 1-6). IEEE.
15. Jabed, M. M. I., Khawer, A. S., Ferdous, S., Niton, D. H., Gupta, A. B., & Hossain, M. S. (2023). Integrating Business Intelligence with AI-Driven Machine Learning for Next-Generation Intrusion Detection Systems. International Journal of Research and Applied Innovations, 6(6), 9834-9849.
16. Kumbum, P. K., Adari, V. K., Chunduru, V. K., Gonepally, S., & Amuda, K. K. (2020). Artificial intelligence using TOPSIS method. Journal of Computer Science Applications and Information Technology, 5(1), 1–7. https://doi.org/10.15226/2474-9257/5/1/00147
17. Gosangi, S. R. (2023). Transforming Government Financial Infrastructure: A Scalable ERP Approach for the Digital Age. International Journal of Humanities and Information Technology, 5(01), 9-15.
18. Cherukuri, B. R. (2024). Serverless computing: How to build and deploy applications without managing infrastructure. World Journal of Advanced Engineering Technology and Sciences, 11(2).
19. "Quantum AI (QAI): Harnessing quantum computing for AI (2024 update)". (2024). PostQuantum.com. Retrieved from https://postquantum.com/quantum-ai-qai/
20. Manda, P. (2023). Migrating Oracle Databases to the Cloud: Best Practices for Performance, Uptime, and Risk Mitigation. International Journal of Humanities and Information Technology, 5(02), 1-7.
21. Joseph, Jimmy. (2024). AI-Driven Synthetic Biology and Drug Manufacturing Optimization. International Journal of Innovative Research in Computer and Communication Engineering. 12. 1138. 10.15680/IJIRCCE.2024.1202069. https://www.researchgate.net/publication/394614673_AIDriven_Synthetic_Biology_and_Drug_Manufacturing_Optimization
22. "Intelligent network optimisation in cloud environments with generative AI and LLMs [v1]". (2024). *Preprints.org*.