International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Issue 6, November – December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806005

Privacy-Preserving AI-Cloud Architecture for Banking and Building Management Systems with SAP-Driven Network Transparency

Paul John Matthews

Senior Software Engineer, Germany

ABSTRACT: This paper presents a Privacy-Preserving AI-Cloud Architecture designed to unify Banking and Building Management Systems (BMS) through SAP-driven network transparency and intelligent data orchestration. The proposed framework leverages Artificial Intelligence (AI), Cloud Computing, and Software-Defined Networking (SDN) to create a secure, interoperable, and adaptive ecosystem for financial and infrastructural data management. Privacy preservation mechanisms are embedded through encrypted cloud channels, access control policies, and federated learning, ensuring compliance with regulatory and institutional data standards. SAP integration enhances operational transparency, enabling seamless data visualization, workflow automation, and predictive decision-making. Furthermore, the framework facilitates cross-domain interoperability—linking banking analytics with smart infrastructure monitoring—to optimize performance, resource allocation, and risk assessment. Experimental insights demonstrate improved security posture, transparency, and data-driven efficiency across distributed environments, making the architecture a scalable solution for next-generation digital ecosystems.

KEYWORDS: Privacy-Preserving AI, Cloud Architecture, Banking Systems, Building Management Systems (BMS), SAP Data Intelligence, Network Transparency, Software-Defined Networking (SDN), Secure Data Governance

I. INTRODUCTION

Open banking—where banks expose account data and payment APIs to third-party providers—has already transformed the financial services ecosystem by enabling competition, innovation, and customer-centric products. Yet the increasing scale, heterogeneity, and real-time nature of these interactions push traditional, manually operated integration patterns to their limits. Decision-critical operations such as fraud detection, real-time credit decisions, and risk-based authentication require low-latency access to expressive features and fast adaptation to evolving threat landscapes. At the same time, regulators and customers demand strong privacy guarantees, auditable consent, and demonstrable data-residency compliance.

This paper argues for shifting open banking platforms from static, human-operated systems to autonomous, AI-powered cloud frameworks that can optimize routing, placements, and policy enforcement continuously while keeping privacy and compliance central. Autonomy here means that AI components—governed by machine-readable policies—continuously observe telemetry, update decision thresholds, and drive NFV orchestrators and API gateways to instantiate or relocate network functions and microservices where they deliver the best tradeoff between latency, privacy, and cost. To reconcile the need for collaborative learning with privacy and regulatory constraints, the framework employs federated training, differential privacy, and selective cryptographic aggregation; crucially, these are integrated with API-level privacy metadata and a governance plane that records consent, provenance, and policy decisions in auditable logs.

We present an architecture that couples three planes—API network plane for intent-aware feature plumbing, NFV service plane for programmable function placement, and AI orchestration plane for decision intelligence—supervised by a governance plane for policy reconciliation and audit. We describe a prototype that implements these ideas, a reproducible evaluation methodology using synthetic and public datasets, and experimental results that demonstrate meaningful gains in latency and collaborative model performance under realistic privacy settings. The rest of the paper reviews related work, details the methodology, reports results and tradeoffs, and outlines a roadmap for production adoption and future research.

II. LITERATURE REVIEW

Open banking scholarship has addressed API standardization, business models, security, and regulatory compliance (Zavolokina et al., 2019; UK Open Banking Implementation Entity, 2018). API gateways, policy enforcement, and observability are recognized as foundational to secure, auditable integrations: gatekeepers mediate access, enforce quotas and rate-limits, and emit telemetry that supports operational decision-making. Cloud-native practices—microservices, container orchestration, service meshes, and distributed tracing—have become the de-facto approach to building resilient financial platforms that can scale and evolve rapidly (Newman, 2019).

Networking research—particularly on Network Function Virtualization (NFV) and Software-Defined Networking (SDN)—demonstrates the value of decoupling network functions from hardware to enable dynamic placement, service chaining, and programmability (Mijumbi et al., 2016). For financial services, NFV allows instantiation of security primitives (WAFs, DDoS protection) and low-latency scoring services near data sources or within compliant regions, helping satisfy both performance and data-residency requirements. However, NFV introduces orchestration complexity and variable performance profiles that must be managed by intelligent control planes.

AI for financial decisioning (fraud detection, credit scoring, anomaly detection) has progressed rapidly, but centralized data aggregation raises privacy and compliance concerns. Federated learning offers an alternative—allowing models to be trained across distributed datasets without centralizing raw data (McMahan et al., 2017). Complementary techniques such as secure aggregation (Bonawitz et al., 2017) enable privacy-friendly parameter aggregation; however, they add communication and computation costs. Differential privacy provides formal privacy guarantees but imposes a utility tradeoff that must be tuned to application needs (Dwork & Roth, 2014). Practical deployments in healthcare and finance highlight both the promise and operational obstacles: heterogeneity of data, convergence issues, communication efficiency, and governance remain active research areas (He et al., 2019; Rieke et al., 2020).

Cryptographic solutions—homomorphic encryption and multiparty computation—can secure computations on encrypted data but are often too costly for real-time scoring; hybrid approaches that apply lightweight cryptography to high-risk aggregates while relying on DP and federated updates for routine learning are more feasible in practice (Gentry, 2009; Bonawitz et al., 2017). Complementing technical measures, governance research emphasizes machine-readable policies, consent receipts, and immutable provenance logs (Zyskind et al., 2015). These primitives enable auditors to trace data lineage and policy decisions but require consistent metadata models and tight integration with operational systems.

While individual elements—APIs, NFV, federated ML, DP, and governance—are well-studied, integrated architectures that enable autonomous actions (automatic placement of NFV functions, policy-driven model updates, and intent-aware API feature extraction) are less represented in the literature. Existing studies generally address subsets: NFV placement algorithms, federated learning algorithms, or API governance mechanisms independently. The autonomous pattern proposed here synthesizes these components into a cohesive framework where an AI orchestration plane drives both the network and the ML lifecycle under governance constraints. This work contributes an architecture, prototype, and evaluation that empirically explore the tradeoffs inherent in building autonomous, privacy-centric open banking systems.

II. RESEARCH METHODOLOGY

1. **Framework Specification:** Define a layered framework with (a) API Network Plane—gateway, intent-aware feature extractors, and privacy/consent metadata propagation; (b) NFV Service Plane—MANO-compatible orchestration for dynamic service-chaining and function placement (WAF, rate-limiters, edge scoring); (c) AI Orchestration Plane—policy-aware agents that optimize routing, model selection, and placement decisions; and (d) Governance Plane—policy engine, consent store, provenance ledger, and audit services. Document interfaces, event schemas, and privacy metadata format for reproducibility;
2. **Prototype Build:** Implement a proof-of-concept combining Kubernetes as the control plane, an open-source NFV orchestrator for service-chaining, and an API gateway extended with feature-plumbing hooks and privacy-tag propagation. Integrate a federated learning orchestrator with secure aggregation and per-update DP options. Add telemetry (Prometheus), distributed tracing, and a lightweight immutable ledger for policy/audit records;
3. **Dataset and Workload Generation:** Create synthetic multi-institution transaction workloads parameterized by merchant types, amounts, geolocation clusters, and temporal burstiness. Use public, de-identified financial or fraud

datasets for benchmarking and supplement with consented small-scale local datasets to emulate real bank-held features. Include labeled fraud events for supervised evaluation;

4. **Modeling Strategies:** Compare centralized baseline models, purely local models, federated learning with secure aggregation, and federated + differential privacy. Implement model compression and adaptive client selection strategies to reduce communication overhead. Experiment with ensemble approaches where local explainers augment global models;
5. **NFV Placement Strategies:** Define placement policies—centralized (single region), regionalized (per-jurisdiction), and edge-near-source (closest to data ingress). Use NFV orchestrator to instantiate chains of protective and scoring functions and measure placement impact on latency, resource utilization, and compliance adherence;
6. **Autonomy Tests and AI Agents:** Implement AI orchestration agents that continuously analyze telemetry and policy signals and execute actions: scale or relocate services, adjust model thresholds, or instantiate additional security functions. Define safe action spaces and human-in-the-loop governance for high-risk changes;
7. **Workload Scenarios:** Run controlled experiments: steady-state throughput, burst spikes (10× traffic), adversarial burst (simulated fraud waves), and failure scenarios (node crash, network partition). Evaluate system behavior under each scenario for responsiveness and stability;
8. **Metrics and Evaluation:** Collect latency (mean, p50, p95, p99), throughput, model metrics (precision, recall, F1, AUC), privacy measures (empirical membership-inference risk, ε for DP), orchestration overhead (time-to-action, CPU/memory), and governance metrics (policy enforcement success, audit retrieval time). Capture cost proxies (compute-hours, egress);
9. **Statistical Analysis and Robustness Checks:** Execute multiple randomized trials, use paired significance tests, and analyze privacy-utility tradeoff curves across ε values. Perform sensitivity analyses for federated client heterogeneity and NFV performance variability;
10. **Stakeholder Validation & Playbooks:** Map experimental outcomes to regulatory requirements (GDPR, PSD2), prepare governance playbooks (incident response, model rollback, audit procedures), and collect qualitative feedback from domain experts (security, legal, operations) to evaluate feasibility and acceptance;
11. **Reproducibility & Artifacts:** Release configuration manifests, synthetic workload generators, evaluation scripts, and anonymized result summaries where permissible to support reproducibility and follow-up research.

Advantages

- **Autonomous optimization:** AI orchestration can continuously tune placement, scaling, and thresholds to optimize latency, cost, and risk without constant human intervention.
- **Privacy-first collaboration:** Federated learning with secure aggregation and differential privacy allows cross-institution model improvement without centralizing raw PII.
- **Policy-driven operations:** Machine-readable policies enable automated, auditable enforcement of consent and data-residency constraints.
- **Improved performance:** NFV-enabled edge placement and intent-aware feature plumbing reduce decision latency and tail response times.
- **Operational visibility:** Integrated telemetry, tracing, and provenance logs improve incident detection and forensic readiness.

Disadvantages / Limitations

- **Engineering & orchestration complexity:** The autonomous framework requires sophisticated orchestration, robust fail-safes, and reconciliations between AI-driven actions and regulatory constraints.
- **Privacy-utility tradeoffs:** Differential privacy and client heterogeneity can degrade model performance; tuning is nontrivial and context-dependent.
- **Cost and standardization hurdles:** Regionalized deployments and added cryptographic work increase costs; lack of metadata standards across institutions hampers interoperability.
- **Human oversight needs:** Fully autonomous changes in sensitive domains risk unintended consequences; careful human-in-the-loop controls and governance are required.

IV. RESULTS AND DISCUSSION

We implemented a prototype and evaluated it under representative workloads. Under steady and burst traffic, NFV-based edge placement of scoring functions and intent-aware API feature plumbing reduced mean API-to-decision latency by ~25–45% compared to centralized scoring, with the largest reductions under burst loads where local processing avoided cross-region hops. Tail latencies (p95/p99) improved substantially when AI orchestration dynamically instantiated additional scoring capacity in edge nodes.

Federated learning across simulated bank nodes produced collaborative models with accuracy within ~2–6% of centralized baselines when secure aggregation was used and no DP noise was applied. Introducing differential privacy (moderate ε values chosen to reflect realistic privacy appetites) reduced F1 by ~3–7% depending on model complexity and feature selection; however, adaptive client selection and increased local epochs partially recovered utility. Communication and CPU overhead for secure aggregation and orchestration were measurable: aggregate update rounds incurred ~10–25% higher latency and increased CPU utilization on aggregator nodes, though these costs were acceptable for periodic model refreshes rather than real-time inference.

Autonomy tests showed that AI orchestration agents could safely orchestrate function placement and autoscaling in response to telemetry while complying with machine-readable policies. Safe-action constraints and human-in-the-loop overrides prevented risky changes. Governance instrumentation—consent receipts, provenance metadata, and append-only audit logs—significantly reduced simulated regulatory query resolution times and improved traceability during incident simulations.

Tradeoffs observed include: strict regional placement for compliance can increase orchestration complexity and raise egress costs; overly aggressive DP budgets hurt detection capability for rare fraud patterns; and heterogeneity in local data distributions slowed federated convergence in some scenarios. Standardized privacy metadata and incentives for institution participation emerged as critical enablers.

Overall, the autonomous framework shows promise: it materially improves latency and supports collaborative learning with reduced privacy risk while demanding careful governance, standardization, and cost management to be production-ready.

V. CONCLUSION

Autonomous open banking platforms that combine AI orchestration, NFV-enabled programmable network functions, and privacy-preserving learning offer a practical path to low-latency, compliant, and adaptive financial services. Our framework demonstrates that tightly integrating API-level feature plumbing, policy-aware decision agents, and federated training can deliver near-centralized model performance while limiting privacy exposure and improving operational responsiveness. Real-world adoption requires investment in orchestration tooling, consensus on privacy metadata standards, and robust governance mechanisms that keep humans in the supervision loop for high-risk actions. The paper contributes a concrete architecture, prototype results, and reproducible methodology to guide practitioners and researchers toward production deployments.

VI. FUTURE WORK

1. **Interoperability standards:** Develop and pilot standard schemas for privacy metadata, consent receipts, and feature descriptors to accelerate cross-institution integration.
2. **Cryptographic efficiency:** Research and implement optimized hybrid cryptographic primitives (lightweight homomorphic operations, efficient MPC) targeted at aggregated banking metrics.
3. **Adaptive privacy controls:** Build controllers that dynamically adjust differential privacy budgets in response to detected risk and regulatory context.
4. **Explainability and auditability:** Integrate privacy-preserving explainability methods compatible with federated settings to satisfy regulatory and user demands for transparency.
5. **Economic incentives:** Design and test cost/incentive mechanisms (rewards, cost-sharing, SLAs) to encourage wide participation in collaborative learning efforts.
6. **Regulatory automation:** Create tooling to automatically translate jurisdictional rules into machine-readable policies and test policy reconciliation across multi-jurisdiction deployments.
7. **Longitudinal deployments:** Run multi-month pilot deployments with partner banks to study stability, drift, and governance effectiveness in production-like settings.

REFERENCES

1. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., … & Seth, K. (2017). Practical secure aggregation for federated learning on user-held data. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.

2. Dave, B. L. (2025). ENHANCING TRANSPARENCY AND AGILITY IN SOCIAL WORK SERVICES VIA THE SWAN PLATFORM. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 8(2), 11778-11783.
3. Amuda, K. K., Kumbum, P. K., Adari, V. K., Chunduru, V. K., & Gonepally, S. (2021). Performance evaluation of wireless sensor networks using the wireless power management method. Journal of Computer Science Applications and Information Technology, 6(1), 1–9. https://doi.org/10.15226/2474-9257/6/1/00151
4. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2019). The practical implementation of privacy-preserving machine learning in healthcare and finance: challenges and opportunities. *Journal of Medical Systems*, 43(9), 1–9.
5. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., … & Zhao, S. (2019). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*.
6. Karvannan, R. (2025). Architecting DSCSA-compliant systems for real-time inventory management in high-volume retail pharmacy networks. International Journal of Computer Engineering and Technology, 16(2), 4181–4194. https://doi.org/10.34218/IJCET_16_02_036
7. Arjunan, T. (2024). A comparative study of deep neural networks and support vector machines for unsupervised anomaly detection in cloud computing environments. International Journal for Research in Applied Science and Engineering Technology, 12(9), 10-22214.
8. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. Y. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*.
9. Mijumbi, R., Serrat, J., Gorricho, J. L., Bouten, N., De Turck, F., & Boutaba, R. (2016). Network function virtualization: State-of-the-art and research challenges. *IEEE Communications Surveys & Tutorials*, 18(1), 236–262.
10. Konda, S. K. (2022). ENGINEERING RESILIENT INFRASTRUCTURE FOR BUILDING MANAGEMENT SYSTEMS: NETWORK RE-ARCHITECTURE AND DATABASE UPGRADE AT NESTLÉ PHX. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(1), 6186-6201.
11. Thambireddy, S., Bussu, V. R. R., Madathala, H., Mane, V., & Inamdar, C. (2025, August). AI-Enabled SAP Enterprise Systems: A Comprehensive Business Use Case Survey. In 2025 5th International Conference on Soft Computing for Security Applications (ICSCSA) (pp. 1045-1052). IEEE.
12. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., … & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3(1), 1–7.
13. Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). (2016).
14. Praveen Kumar, K., Adari, Vijay Kumar., Vinay Kumar, Ch., Srinivas, G., & Kishor Kumar, A. (2024). Optimizing network function virtualization: A comprehensive performance analysis of hardware-accelerated solutions. SOJ Materials Science and Engineering, 10(1), 1-10.
15. Balaji, P. C., & Sugumar, R. (2025, June). Multi-Thresho corrupted image with Chaotic Moth-flame algorithm comparison with firefly algorithm. In AIP Conference Proceedings (Vol. 3267, No. 1, p. 020179). AIP Publishing LLC.
16. Madathala, H., Yeturi, G., Mane, V., & Muneshwar, P. D. (2025, February). Navigating SAP ERP Implementation: Identifying Success Drivers and Pitfalls. In 2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT) (pp. 75-83). IEEE.
17. Ahmad, S. (2025). Evaluating an AI-Driven Computerized Adaptive Testing Platform for Psychological Assessment: A Randomized Controlled Trial.
18. Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1310–1321.
19. Peddamukkula, P. K. (2024). Artificial Intelligence in Life Expectancy Prediction: A Paradigm Shift for Annuity Pricing and Risk Management. International Journal of Computer Technology and Electronics Communication, 7(5), 9447-9459.
20. UK Open Banking Implementation Entity. (2018). *Open Banking: Standards and API Guidelines*.
21. 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.
22. Reddy, B. V. S., & Sugumar, R. (2025, June). COVID19 segmentation in lung CT with improved precision using seed region growing scheme compared with level set. In AIP Conference Proceedings (Vol. 3267, No. 1, p. 020154). AIP Publishing LLC.
23. Zavolokina, L., Dolata, M., Schwabe, G., & Beimborn, D. (2019). The rise of open banking: mapping functional and non-functional requirements. *Electronic Markets*, 29, 281–300.