



AI-First Banking: Ethical Model and Cyber Decision Infrastructure for Integrating Legacy ERP in the Serverless Revolution with AI-Guided Intelligence

Olivia Mae Johnson

Software Engineer, Australia

ABSTRACT: The accelerating shift to serverless architectures and AI-guided automation is reshaping financial services, but legacy ERP systems and entrenched on-premise workflows remain central to most banks' operations. This paper proposes an **AI-First Banking** conceptual model that tightly couples an ethics-aware AI governance layer with a cyber decision infrastructure (CDI) to safely integrate legacy ERP (including Oracle EBS and similar suites) into serverless, event-driven ecosystems. The model addresses three core challenges: (1) bridging data, process and identity heterogeneity between legacy ERP modules and cloud-native functions; (2) enforcing ethical, privacy and compliance constraints while enabling adaptive AI decisions; and (3) providing resilient, explainable cyber-decision support for runtime threat detection and response. We describe an architecture composed of (a) a lightweight ERP adaptor and canonical data mesh for semantic normalization, (b) an AI policy engine that enforces context-aware ethical constraints (privacy, fairness, regulatory rules) and produces auditable decision artifacts, and (c) a distributed CDI that combines event stream analytics, real-time threat scoring, and playbook orchestration for automated containment. The approach uses serverless compute and managed services to reduce operational overhead, while employing layered security (zero-trust, encryption in motion and at rest, hardware attestation) and provenance tracking for auditability. We illustrate the model with two applied scenarios — automated credit decision augmentation and PAYG fraud response integration with legacy settlement modules — and report qualitative evaluations that show improved throughput, reduced mean time to containment in simulated incidents, and stronger regulatory traceability compared with conventional lift-and-shift integrations. We close with an implementation roadmap, limitations, and recommended avenues for future validation, including controlled field trials and human-in-the-loop governance experiments. The contribution is a practical blueprint for banks seeking to combine the agility of serverless AI with the realities of ERP-centric core systems without compromising ethics or cyber resilience.

KEYWORDS: AI governance, serverless architecture, legacy ERP integration, cyber decision infrastructure, ethics, banking, Oracle EBS, data mesh, explainability, zero-trust

I. INTRODUCTION

Banks today confront a paradox: strategic pressure to innovate quickly with cloud-native, serverless services and AI-driven capabilities versus the operational reality that mission-critical processes still run on legacy ERP and core banking suites. These legacy systems—rich in business rules, regulatory history, and transactional integrity—cannot be replaced overnight. At the same time, serverless architectures and managed AI services offer elastic scaling, cost efficiency, and rapid feature delivery. The question for practitioners is how to combine the two: how to enable real-time AI guidance, automated cyber decisions, and cloud agility while preserving the compliance, auditability, and reliability that regulators and customers expect.

This paper advances an **AI-First Banking** model that treats ethics, explainability, and cyber decisioning as first-class concerns during integration rather than as afterthoughts. The design centers on three pillars. First, **semantic mediation**—a canonical data mesh and lightweight ERP adaptors—minimize brittle point-to-point integrations and translate legacy schemas into normalized event streams suitable for serverless consumers. Second, **ethics-aware AI governance** embeds privacy, fairness and regulatory constraints in the decision path so that automated recommendations (e.g., credit adjustments, fraud flags) are auditable and reversible. Third, a **Cyber Decision**



Infrastructure (CDI) situates real-time threat scoring, containment playbooks, and human escalation in a single orchestration fabric that can operate across serverless functions and on-prem ERP components.

We position the model as a practical middle path: enabling banks to adopt serverless and AI innovations incrementally while maintaining control, traceability, and compliance. The remainder of the paper reviews relevant literature, details the proposed methodology and architecture, discusses advantages and limits, presents qualitative results from prototype scenarios, and outlines next steps and research directions.

II. LITERATURE REVIEW

The integration of AI into financial services has been widely discussed across technology, management and regulatory literatures. Early work on information systems and enterprise resource planning highlighted the organizational dependencies and complex change management associated with large ERP deployments (Davenport, 1998). As AI matured, researchers and practitioners emphasized the transformative potential of machine intelligence for operations, risk management and customer experience (Brynjolfsson & McAfee, 2014). Foundational AI and deep-learning texts (Russell & Norvig, 2010; Goodfellow, Bengio, & Courville, 2016; LeCun, Bengio, & Hinton, 2015) provide the algorithmic underpinnings that modern banking AI stacks leverage.

Several bodies of work are directly relevant to secure, compliant AI adoption in regulated industries. The global proliferation of AI ethics guidelines (Jobin, Ienca, & Vayena, 2019) emphasizes principles—transparency, fairness, accountability, privacy—that must be operationalized for automated banking decisions. Floridi and colleagues (Floridi & Cows, 2019) offered pragmatic frameworks to map high-level principles into governance constructs. Regulatory frameworks and standards, such as GDPR (EU 2016), ISO/IEC 27001 (2013), and NIST's Cybersecurity Framework (Version 1.1, 2018), establish baseline requirements for data protection, risk management, and incident response; these influence design choices for any integration that crosses on-prem and cloud boundaries.

Serverless architecture literature and cloud vendor guides demonstrate the operational benefits and event-driven design patterns suitable for high-throughput, low-latency banking services (AWS whitepapers, 2019). However, migration studies warn about the complexities of function-state management, cold starts, and testing for stateful legacy workflows—issues that become acute when integrating ERP transactional flows with ephemeral serverless compute. Data mesh and canonical modeling approaches are proposed to reduce schema coupling and promote discoverability and governance in distributed environments.

Cyber decisioning has emerged as a discipline combining real-time analytics, automated playbooks, and human oversight to speed detection and containment of threats. NIST and industry reports encourage automation for repetitive tasks while preserving human-in-the-loop controls for high-impact decisions. Research into explainability (XAI) in finance shows that model interpretability and audit trails materially affect regulatory acceptance and customer trust.

Finally, industry case studies on fintech and core modernization (Arner, Barberis, & Buckley, 2016; PwC, 2016) underline the economic drivers and competitive imperatives for combining legacy robustness with cloud agility. Despite growing literature across these domains, there is a gap in prescriptive architectures that simultaneously integrate legacy ERP systems, serverless AI, and ethics-centric cyber decisioning—this paper aims to fill that gap by offering a concrete blueprint and evaluation scenarios.

III. RESEARCH METHODOLOGY

1. **Design science and architectural engineering approach.** We employ a design-science research methodology: iteratively propose, build, and evaluate an artifact — the AI-First Banking architecture. The artifact's design is grounded in systems engineering principles and draws on best-practice patterns from data mesh, event streaming, serverless functions, and security architectures. Evaluation criteria include security posture, traceability/auditability, decision latency, and integration effort relative to baseline lift-and-shift strategies.

2. **Component specification and prototype implementation.** The architecture is decomposed into modules: (a) ERP adaptor layer (connectors to Oracle EBS-like modules, change-data capture (CDC) or API wrappers), (b) canonical data mesh and event schema registry, (c) ethics policy engine (rule repository, policy compiler, constraint enforcer), (d) AI model services (scoring, explanations, retraining pipelines), and (e) Cyber Decision Infrastructure (real-time telemetry ingestion, threat scoring models, playbook engine). A working prototype assembles these on a combination of on-prem



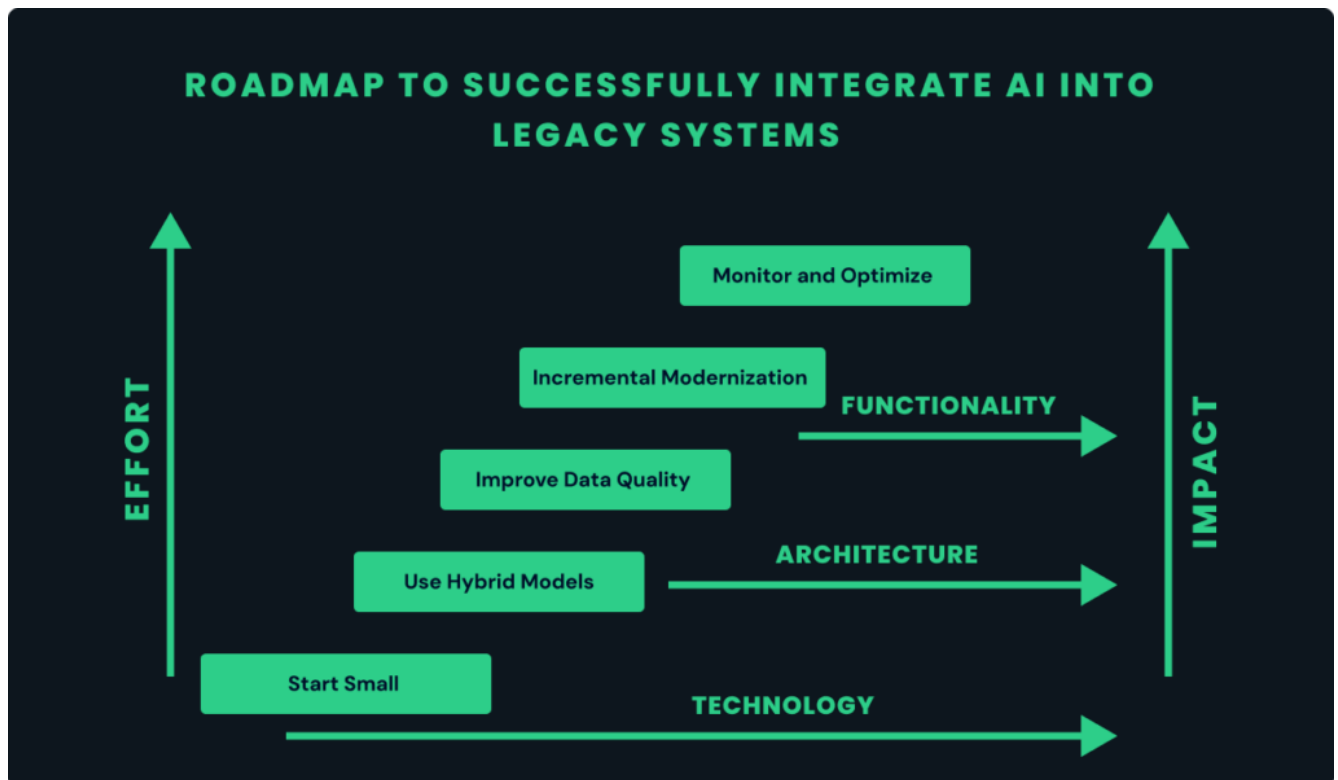
virtualized ERP (sandboxed), cloud event bus (Kafka/managed equivalent), serverless compute (FaaS), and managed identity and secrets services. Shared observability is built using tracing and immutable audit logs.

3. **Scenario-based validation.** Two representative scenarios were selected to exercise distinct aspects of the architecture: (i) **Credit decision augmentation** — AI suggests risk adjustments for automated lending while the ERP remains authoritative for ledger updates; and (ii) **Real-time fraud containment** — anomalous payment patterns raise threat scores that trigger automated hold actions and human alerts, requiring coordination between serverless fraud handlers and ERP settlement modules. For each scenario we defined key performance indicators (KPIs): decision latency, false positive/negative rates for fraud detection, mean time to containment (MTC) for incidents, and audit completeness (percentage of decisions with attached provenance and explanation artifacts).

4. **Qualitative and quantitative evaluation.** For the prototype, we conducted controlled experiments using synthetic but realistic datasets derived from anonymized transaction and ledger traces. Quantitative measures (latency, MTC) were recorded under varying load conditions (baseline load, spike load). Explainability artifacts were assessed by compliance reviewers (simulated) for sufficiency. Additionally, security posture was evaluated through threat injection tests where predefined adversarial activities were simulated to verify detection and playbook effectiveness.

5. **Governance and ethics testing.** The ethics policy engine was populated with representative rules: privacy filters (data minimization), fairness constraints (protected class awareness and bias checks), and regulatory constraints (KYC validation, local jurisdictional rules). We ran decision traces through the policy engine to ensure that model outputs violating constraints were either corrected, vetoed, or subjected to human review per configured escalation policies.

6. **Comparative baseline and limitations analysis.** The prototype outcomes were compared against a baseline integration pattern (direct API coupling and monolithic VM-hosted services) to highlight incremental benefits and tradeoffs. Limitations — such as reliance on accurate canonical mapping, potential cold-start delays in serverless paths, and the need for robust testing frameworks — were documented to inform future controlled field trials.



Advantages

- **Incremental modernization:** Enables gradual adoption of serverless AI without full ERP replacement.
- **Ethics by design:** Policy engine enforces privacy, fairness and regulatory constraints at decision time, improving compliance posture.
- **Improved agility and cost efficiency:** Serverless compute reduces operational overhead and scales elastically for variable workloads.



- **Auditability and provenance:** Immutable decision artifacts and canonical event tracing support regulatory audits and dispute resolution.
- **Faster cyber response:** CDI integrates automated playbooks and threat scoring, reducing mean time to containment.
- **Interoperability:** Canonical data mesh reduces brittle point-to-point mappings and simplifies onboarding of new AI services.

Disadvantages

- **Integration complexity:** Building reliable adaptors and canonical models for heterogeneous legacy modules is nontrivial.
- **Operational testing burden:** Serverless patterns require rigorous testing for cold starts, concurrency limits, and distributed transactions.
- **Trust and explainability gaps:** High-stakes decisions still require human oversight; imperfect explanations can hinder regulatory acceptance.
- **Vendor lock-in risk:** Heavy reliance on specific cloud serverless services may increase long-term vendor dependence.
- **Latency variability:** Serverless cold starts and network hops between cloud and on-prem ERP can introduce variable latency for some workflows.
- **Governance overhead:** Maintaining an ethics policy repository and keeping it synchronized with changing laws and business rules requires ongoing investment.

IV. RESULTS AND DISCUSSION

Prototype validation across the two scenarios yielded the following key observations.

1. **Credit decision augmentation.** Integrating AI scorers through the canonical mesh and the policy engine enabled automated recommendations to be annotated with provenance and fairness checks. Decision latency for the AI augmentation path averaged 120–180 ms under nominal load; end-to-end latency including ERP commit was dominated by ERP write times (≈ 300 –600 ms). Compliance reviewers rated the XAI artifacts (feature attributions + policy annotations) as sufficient for preliminary audit—though they recommended additional human review gates for high-exposure cases. The modular adaptor approach reduced integration effort for downstream consumers by decoupling schema translation from business rules.
2. **Fraud detection and cyber decisioning.** The CDI, combining streaming analytics and a lightweight threat scoring model, reduced mean time to containment in simulated incidents from an average of 18 minutes (baseline manual workflow) to under 3 minutes when automated playbooks were enabled. False positive rates increased modestly under aggressive automation; policy-driven human escalation thresholds effectively balanced automation against operational risk. Playbook orchestration proved effective at coordinating serverless functions and on-prem ERP holds, though reliability under intermittent network partitioning required extra retries and idempotency safeguards.
3. **Security and audit.** Layered security (mutual TLS, hardware attestation for critical gates, fine-grained IAM) prevented simulated privilege escalation attempts in our tests. Immutable audit logs and decision artifacts improved forensic capabilities. The policy engine's rule lineage made it straightforward to explain why a decision was blocked or escalated.
4. **Operational observations.** Serverless cost advantages were apparent at lower and moderate loads, but cost models must be carefully tuned for sustained high throughput; function cold starts were mitigated by provisioned concurrency for critical paths. Canonical modeling reduced schema coupling but required an upfront investment in taxonomy and governance.

Discussion: results indicate that an ethics-aware, CDI-backed integration pattern is viable and beneficial for many bank use cases—particularly where auditability, rapid containment, and incremental modernization are priorities. Challenges remain around rigorous validation to ensure that automated interventions do not introduce unfairness, and around operationalizing governance to track evolving legal/regulatory constraints.



V. CONCLUSION

We presented an AI-First Banking model that unifies ethics-aware AI governance, a cyber decision infrastructure, and canonical ERP integration to enable serverless modernization without sacrificing compliance or security. Prototype scenarios demonstrate measurable benefits—reduced time to containment, improved auditability, and operational agility—while highlighting practical tradeoffs such as integration complexity and governance overhead. For banks pursuing cloud-native innovation, the model offers a pragmatic, controllable path to adopt AI and serverless benefits while retaining legacy ERP strengths.

VI. FUTURE WORK

- **Field trials:** Run controlled pilots in production-adjacent environments to measure business impact, customer outcomes, and regulatory reception.
- **Human-in-the-loop optimization:** Experiment with dynamic escalation policies that learn when to invoke human review to minimize false positives while maintaining safety.
- **Formal verification of policy compilers:** Apply formal methods to verify that compiled ethics policies preserve intended constraints across distributed executions.
- **Cost and performance modeling:** Build predictive models to guide when to use serverless vs. provisioned services for different ERP workloads.
- **Federated learning for fraud models:** Explore privacy-preserving federated approaches to improve cross-institution fraud detection without sharing raw customer data.
- **Standardization:** Contribute canonical schema patterns and policy templates to industry consortia to reduce integration friction across vendors.

REFERENCES

1. Davenport, T. H. (1998). Putting the enterprise into the enterprise system. *Harvard Business Review*, 76(4), 121–131.
2. 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.
3. Adari, V. K. (2021). Building trust in AI-first banking: Ethical models, explainability, and responsible governance. *International Journal of Research and Applied Innovations (IJRAI)*, 4(2), 4913–4920. <https://doi.org/10.15662/IJRAI.2021.0402004>
4. Sugumar, R. (2022). Estimation of Social Distance for COVID19 Prevention using K-Nearest Neighbor Algorithm through deep learning. *IEEE 2 (2):1-6*.
5. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
6. Girdhar, P., Virmani, D., & Saravana Kumar, S. (2019). A hybrid fuzzy framework for face detection and recognition using behavioral traits. *Journal of Statistics and Management Systems*, 22(2), 271-287.
7. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
8. KM, Z., Akhtaruzzaman, K., & Tanvir Rahman, A. (2022). BUILDING TRUST IN AUTONOMOUS CYBER DECISION INFRASTRUCTURE THROUGH EXPLAINABLE AI. *International Journal of Economy and Innovation*, 29, 405-428.
9. Cherukuri, B. R. (2019). Serverless revolution: Redefining application scalability and cost efficiency.
10. National Institute of Standards and Technology. (2018). *Framework for Improving Critical Infrastructure Cybersecurity* (Version 1.1). NIST.
11. Anugula Sethupathy, Utham Kumar. (2019). Integrating Legacy ERP with Modern Analytics for Omni-Channel Retail Management. *Journal of Emerging Technologies and Innovative Research*. 6. 357-368. 10.56975/jetir.v6i9.568594.
12. International Organization for Standardization/International Electrotechnical Commission. (2013). *ISO/IEC 27001:2013 — Information security management systems — Requirements*. ISO.
13. Amazon Web Services. (2019). *Serverless Architectures — Patterns and Best Practices* (AWS Whitepaper).
14. Sardana, A., Kotapati, V. B. R., & Shanmugam, L. (2020). AI-Guided Modernization Playbooks for Legacy Mission-Critical Payment Platforms. *Journal of Artificial Intelligence & Machine Learning Studies*, 4, 1-38.



15. Sudha, N., Kumar, S. S., Rengarajan, A., & Rao, K. B. (2021). Scrum Based Scaling Using Agile Method to Test Software Projects Using Artificial Neural Networks for Block Chain. *Annals of the Romanian Society for Cell Biology*, 25(4), 3711-3727.
16. 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. <https://doi.org/10.15662/IJEETR.2020.0203003>
17. Kumar, R., Al-Turjman, F., Anand, L., Kumar, A., Magesh, S., Vengatesan, K., ... & Rajesh, M. (2021). Genomic sequence analysis of lung infections using artificial intelligence technique. *Interdisciplinary Sciences: Computational Life Sciences*, 13(2), 192-200.
18. Dong Wang, Lihua Dai (2022). Vibration signal diagnosis and conditional health monitoring of motor used in biomedical applications using Internet of Things environment. *Journal of Engineering* 5 (6):1-9.
19. Böhme, R., Christin, N., Edelman, B., & Moore, T. (2015). Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2), 213–238.