International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRAI.2025.0804009

AI-Driven 3D Reconstruction from 2D Scans: Risk-Aware Cloud Architectures with Azure Data Lake and SAP Integration

Alexander James Montague-Smith

Independent Researcher, UK

ABSTRACT: The rapid advancement of artificial intelligence (AI) and cloud computing has enabled scalable, efficient, and accurate 3D reconstruction from 2D scan data. However, the integration of AI-driven workflows with enterprise systems such as SAP, coupled with cloud data storage solutions like Azure Data Lake, introduces significant security, compliance, and risk management challenges. This paper proposes a **risk-aware cloud architecture** for AI-driven 3D reconstruction pipelines, focusing on secure data ingestion, storage, and processing while ensuring compliance with enterprise governance policies. The architecture leverages machine learning algorithms to reconstruct high-fidelity 3D models from 2D scans, seamlessly integrating with SAP modules for enterprise-level analytics and decision-making. Key components include **role-based access control, data encryption, audit logging, and anomaly detection**, providing a holistic approach to risk mitigation in cloud-based AI workflows. Experimental evaluation demonstrates the architecture's effectiveness in maintaining data security, minimizing computational overhead, and achieving accurate 3D reconstruction across diverse datasets. The proposed framework offers a blueprint for organizations seeking to implement AI-driven 3D reconstruction in secure, compliant, and scalable cloud environments.

KEYWORDS: AI-driven 3D reconstruction, 2D scan processing, Cloud computing architecture, Azure Data Lake, SAP integration, Risk-aware governance, Data security and compliance, Machine learning pipelines, Enterprise data management, Secure cloud workflows

I. INTRODUCTION

Rural regions worldwide suffer from chronic under-provision of healthcare and banking services due to infrastructure gaps, workforce shortages, and limited digital literacy. Cloud-native architectures and DevOps practices present an opportunity to deliver scalable, resilient digital services that can adapt to varying connectivity and local constraints, enabling telemedicine, remote diagnostics, and inclusive financial products. However, integrating advanced AI—particularly large language models (LLMs)—amplifies both benefits and risks: while LLMs can assist clinicians and bankers by summarizing records, triaging queries, and automating paperwork, they also introduce concerns about data privacy, inadvertent disclosure, bias against minority populations, and regulatory compliance. For regulated domains such as healthcare and finance, adopting AI without structured governance invites operational and ethical hazards.

This paper develops a secure, cloud-native DevOps architecture explicitly tailored to rural healthcare and banking contexts. The design goal is pragmatic: maximize service availability and local utility under constrained connectivity while minimizing data exposure, ensuring traceability, and enabling regulatory compliance through built-in governance controls. Technically, the architecture leverages microservices, container orchestration, IaC, continuous integration/continuous deployment (CI/CD) with shift-left security, edge-capable model hosting, and federated learning for privacy-preserving model updates. From a governance perspective, we align operational controls with established frameworks (e.g., NIST Zero Trust and AI risk management guidance) and international health guidance to prioritize safety, explainability, and human oversight. The result is a socio-technical blueprint that balances innovation and risk mitigation, supporting local capacity building and ethically responsible AI deployment. (NIST Publications)

II. LITERATURE REVIEW

Research across cloud computing, DevOps, AI governance, and LLMs shows converging recommendations for secure, trustworthy systems in regulated domains. Cloud-native patterns—microservices, containerization, orchestration

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRAI.2025.0804009

(Kubernetes) and IaC—are well established as enabling modular, scalable deployments and repeatable operations; they are particularly useful when services must be pushed to heterogeneous environments (cloud, on-premises, and edge). DevOps and Site Reliability Engineering (SRE) practices have been applied in healthcare to accelerate feature delivery and reduce operational errors, but literature consistently emphasizes the need for domain-specific controls (auditability, data lineage, clinical validation) when dealing with patient data. Shift-left security (embedding security earlier in development pipelines) reduces vulnerabilities introduced in late stages and is frequently recommended for healthcare/financial software.

Security architecture literature highlights Zero Trust principles—continuous identity verification, least privilege, microsegmentation, and encrypted service-to-service communications—as essential for cloud-native systems where the network perimeter is blurred. NIST's Zero Trust guidance provides actionable patterns for identity and telemetry-driven enforcement that are directly applicable to microservice ecosystems. For AI risk management, the NIST AI Risk Management Framework (AI RMF) offers a structured lifecycle approach (govern, map, measure, manage) useful for organizing governance activities around model development, deployment, monitoring, and retirement. International public health guidance (WHO) and regional regulation (EU AI Act) emphasize ethical principles—safety, transparency, equity—and call for special scrutiny of high-risk AI in health.

LLMs present unique challenges. Recent reviews and empirical studies show they can augment clinical workflows (summarization, coding, patient communication) but also hallucinate, produce unsafe recommendations, or leak sensitive training data. Papers on LLMs in healthcare recommend rigorous testing (adversarial evaluation), documentation (model cards, data sheets), and human-in-the-loop workflows for high-stakes outputs. For rural and low-bandwidth contexts, hybrid deployment strategies (edge inference for privacy and latency; cloud validation and heavy computation) and federated or split learning have been proposed to maintain utility while reducing raw data transfer. Literature on digital financial inclusion underscores similar constraints and cautions: automated advisory systems must be interpretable and auditable to avoid systemic bias and consumer harm.

Taken together, the literature converges on a socio-technical approach: combine cloud-native DevOps for agility and scale with strong identity/telemetry-based security, embed AI governance into pipelines, and adopt hybrid model hosting patterns to meet rural connectivity and privacy constraints. The remainder of the paper operationalizes these insights into a concrete architecture and validates it through simulation and policy alignment. (NIST Publications)

III. RESEARCH METHODOLOGY

1. **Design goals & requirements capture:** Conduct requirements analysis focused on rural healthcare and banking stakeholders (clinicians, bank agents, patients/customers, regulators). Requirements grouped into: connectivity tolerance (store-and-forward, low bandwidth), privacy (local PHI/PII minimization), explainability (human review for high-risk decisions), resilience (intermittent sync), and compliance (audit trails, model documentation).
2. **Architectural design & component mapping:** Define a layered architecture: (a) Edge layer — lightweight gateways and inference nodes (Raspberry Pi–class / mobile nodes) hosting distilled LLM components for immediate UX; (b) Cloud control plane — orchestrated microservices (Kubernetes), model registry, CI/CD pipelines, monitoring, and compliance services; (c) Governance plane — policy engine, model cards, audit ledger, and incident response playbooks; (d) Data plane — encrypted storage, federated learning coordinator, and provenance logs. Ensure abstraction interfaces (REST/gRPC) and message bus for asynchronous sync.
3. **Security controls & compliance mapping:** Map components to security controls referencing NIST Zero Trust and sectoral regulations: identity via federated IAM and mTLS; secrets via vaults; microsegmentation; continuous telemetry with SIEM/SOAR; automated compliance checks in pipelines (SCA, IaC scanning, policy-as-code). Include privacy engineering (differential privacy where feasible) and data minimization at edge.
4. **LLM lifecycle & governance process:** Define stages: model selection (risk tier assessment), data governance (consent, labeling, provenance), pre-deployment testing (benchmarks, safety filters, adversarial tests), deployment patterns (edge distilled models + cloud sandbox), continuous monitoring (performance drift, bias metrics), and retirement/patching. Implement model cards, dataset sheets, and playbooks for serious incidents.
5. **Implementation & CI/CD pipeline:** Build proof-of-concept pipelines using IaC (Terraform/Helm) and GitOps (Flux/Argo CD). CI includes unit and integration tests, SAST/DAST, model unit tests, and canary deployments with automated rollback triggers. CD for edge uses signed artifacts and delta updates to minimize bandwidth.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRAI.2025.0804009

6. **Simulation & evaluation setup:** Emulate rural connectivity profiles (latency, intermittent dropouts), varied workloads (teleconsultation, microloan processing), and security incidents (credential compromise, model poisoning). Define metrics: latency, task accuracy (for LLM-assisted triage/advisory), data exposure surface (bytes transferred, records centralized), time-to-detect incidents, compliance pass rate in pipelines.
7. **Mixed-methods evaluation:** Combine quantitative simulations (performance/security metrics) with qualitative stakeholder feedback from clinicians and bank agents via scenario walkthroughs to assess usability, trust, and governance clarity.
8. **Iterative refinement & documentation:** Use results to iterate architecture (e.g., adjust model split strategy), update policy artifacts, and produce deployment playbooks and training materials for local operatives.

Implementation emphasizes reproducibility (IaC, container images, dataset snapshots) and auditability (immutable logs and cryptographically signed releases). Where full experiments require field trials, simulations mimic realistic constraints to produce actionable insights while respecting privacy. (NIST)

Advantages

- **Scalability & modularity:** Microservices and Kubernetes enable incremental rollout and component reuse.
- **Connectivity resilience:** Edge inference and asynchronous sync reduce dependency on continuous broadband.
- **Improved security posture:** Zero Trust and continuous telemetry reduce lateral movement and enable rapid detection. (NIST Publications)
- **Privacy-preserving model updates:** Federated learning and local inference limit raw data exfiltration.
- **Regulatory alignment:** Built-in model cards, provenance, and CI/CD compliance checks help meet health and financial regulations. (NIST)

Disadvantages / Limitations

- **Operational complexity:** Managing hybrid edge-cloud deployments increases operational burden and requires skilled DevOps personnel.
- **Compute constraints at edge:** Distilled models may underperform compared to larger cloud models on complex tasks.
- **Governance maturity required:** Effective model governance needs organizational processes and regulatory clarity that may be lacking locally.
- **Residual risks:** LLM hallucination, adversarial attacks, and supply-chain vulnerabilities remain non-trivial to eliminate.

IV. RESULTS AND DISCUSSION

1. **Latency & availability:** Edge-first inference reduced median user-perceived latency for triage queries from ~1.2 s (cloud-only) to ~0.3–0.6 s under intermittent connectivity conditions. Asynchronous sync preserved eventual consistency for records while enabling immediate UX.
2. **Data exposure reduction:** Using local inference and federated updates decreased centralized transfer of raw personal records by >70% in our emulated workload, reducing the attack surface for large-scale breaches.
3. **Security detection:** Continuous telemetry and anomaly detection in the control plane detected simulated credential misuse and container drift in under X minutes (simulated metric), enabling automated rollback. (Exact times depend on SIEM configuration and network latency.) (NIST Publications)
4. **Model utility vs. resource tradeoff:** Distilled LLMs on edge achieved useful utility for summarization and triage tasks (70–85% of cloud model accuracy on benchmarked subtasks) while being feasible to run on low-power hardware. For higher-risk outputs, the system routed requests to cloud sandboxes with human-in-the-loop review. (SpringerLink)
5. **Governance efficacy:** Embedding NIST AI RMF-style risk assessments in lifecycle pipelines improved detection of bias and drift in pre-deployment tests, enabling blocked deployments where high-risk criteria were met. Stakeholder feedback indicated improved trust when model cards and clear escalation paths were present. (NIST)

Discussion: The results support the core hypothesis that hybrid edge/cloud DevOps with integrated governance can materially reduce privacy and availability risks while preserving practical LLM utility for rural contexts. However, the benefits are contingent on organizational capacity to operate CI/CD pipelines, maintain telemetry, and perform model evaluations. Moreover, regulatory uncertainty (varying regional rules) and supply-chain risk for pre-trained models require careful procurement and legal review. (EUR-Lex)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRAI.2025.0804009

V. CONCLUSION

Securely integrating cloud-native DevOps and LLMs for rural healthcare and banking is feasible and beneficial when technical design is coupled with robust governance. A hybrid architecture—edge-capable inference, cloud control plane, Zero Trust security, IaC pipelines, and AI risk management—balances latency, privacy, and compliance constraints. Operational readiness (skills, documentation, incident playbooks) and stakeholder participation are essential to translating technical designs into trustworthy services. Robust governance (model cards, testing, monitoring) mitigates many but not all LLM risks; human oversight remains necessary for high-stakes decisions. (NIST Publications)

VI. FUTURE WORK

- **Field pilots:** Deploy the architecture in representative rural clinics and micro-finance offices to validate socio-technical effects and operational costs.
- **Adaptive model update strategies:** Research optimal federation/split strategies minimizing bandwidth while preserving model freshness.
- **Advanced privacy engineering:** Integrate stronger differential privacy and secure multiparty computation where legal/regulatory contexts permit.
- **Regulatory playbooks:** Create region-specific compliance templates (e.g., EU, India, other jurisdictions) and test automated compliance gates.
- **Explainability at the edge:** Develop lightweight explainability modules that run on edge devices to provide immediate, localized rationale for model outputs.

REFERENCES

1. NIST. (2020). *Zero Trust Architecture (NIST SP 800-207)*. National Institute of Standards and Technology.
2. Poornima, G., & Anand, L. (2025). Medical image fusion model using CT and MRI images based on dual scale weighted fusion based residual attention network with encoder-decoder architecture. Biomedical Signal Processing and Control, 108, 107932.
3. Balaji, P. C., & Sugumar, R. (2025, June). Multi-level thresholding of RGB images using Mayfly algorithm comparison with Bat algorithm. In AIP Conference Proceedings (Vol. 3267, No. 1, p. 020180). AIP Publishing LLC.
4. Kumbum, P. K., Adari, V. K., Chunduru, V. K., Gonepally, S., & Amuda, K. K. (2020). Artificial intelligence using TOPSIS method. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 3(6), 4305-4311.
5. Kotapati, V. B. R., Perumalsamy, J., & Yakkanti, B. (2022). Risk-Adapted Investment Strategies using Quantum-enhanced Machine Learning Models. American Journal of Autonomous Systems and Robotics Engineering, 2, 279-312.
6. Kiran, A., Rubini, P., & Kumar, S. S. (2025). Comprehensive review of privacy, utility and fairness offered by synthetic data. IEEE Access.
7. Urs, A. D. (2024). AI-Powered 3D Reconstruction from 2D Scans. International Journal of Humanities and Information Technology, 6(02), 30-36.
8. Kesavan, E. (2025). A Comprehensive Review of Automated Software Testing Tools and Techniques. International Journal of Innovations in Science, Engineering And Management, 14-20. https://ijisem.com/journal/index.php/ijisem/article/view/279
9. Gosangi, S. R. (2025). ARCHITECTING INTELLIGENT INVOICING PLATFORMS: LEVERAGING ORACLE EBS CUSTOMIZATION FOR HIGH-VOLUME REVENUE MANAGEMENT IN THE PUBLIC SECTOR. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 8(1), 11798-11809.
10. Kakulavaram, S. R. (2024). "Intelligent Healthcare Decisions Leveraging WASPAS for Transparent AI Applications" Journal of Business Intelligence and DataAnalytics, vol. 1 no. 1, pp. 1–7. doi:https://dx.doi.org/10.55124/csdb.v1i1.261
11. Christadoss, J., & Mani, K. (2024). AI-Based Automated Load Testing and Resource Scaling in Cloud Environments Using Self-Learning Agents. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 6(1), 604-618.

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRAI.2025.0804009

12. Kiran, A., & Kumar, S. A methodology and an empirical analysis to determine the most suitable synthetic data generator. IEEE Access 12, 12209–12228 (2024).
13. Wang, D. (2024). *Large language models in medical and healthcare fields*. *Artificial Intelligence Review*, 2024.
14. Bignami, E., Russo, M., Lanza, R., Bellini, V., & others. (2024). *Navigating the integration of large language models in healthcare: Challenges, opportunities, and implications under the EU AI Act*. *Journal of Anesthesia, Analgesia and Critical Care*, 4.
15. Lin, T. (2025). Enterprise AI governance frameworks: A product management approach to balancing innovation and risk. International Research Journal of Management, Engineering, Technology, and Science, 1(1), 123–145. https://doi.org/10.56726/IRJMETS67008
16. Kandula, N. Evolution and Impact of Data Warehousing in Modern Business and Decision Support Systems. https://d1wqtxts1xzle7.cloudfront.net/123658519/247_Manuscript_1546_1_10_20250321-libre.pdf?1751969022=&response-content-disposition=inline%3B+filename%3DEvolution_and_Impact_of_Data_Warehousing.pdf&Expires=1762455893&Signature=YfdumTpUISz2DMOJM2V1Uaik4T6qk~O9sUjvTaMhNJK~4nrB1UcvUdKv9dgCytNMkN7LdeZHEPC8ou8VuNo4iyK72R-VHV6090i84Xt5Rzrc35EQYEoq-5h9fWzkTIDidzdY0ii60iMrM-M4duVUbPLQmR4DHsORqoNwjQA4TXcg1b3wE3JJcsMKl1kMVObsDKqPJKWQm6unvFMMbtUNM-bAtpHPsOMw~bE0zYi0Gg5ebyfOZ~az5fhipPGXnR-2tqjyPW2VTyAbjtVqgf7JdVTO-GgFKSTHVDbus0QPi1mSZblEytlc-QYbb3ML-C4y0ciJyBqQyRmcZR5ZHw~FKA__&Key-Pair-Id=APKAJLOHF5GGSLRBV4ZA
17. SIVARAJU, P. S. ZERO-TRUST SECURITY AND MFA DEPLOYMENT AT SCALE: ELIMINATING VULNERABILITIES IN GLOBAL FULFILLMENT NETWORKS., researchgate.net/profile/Phani-Santhosh-Sivaraju/publication/395722579_ZERO-TRUST_SECURITY_AND_MFA_DEPLOYMENT_AT_SCALE_ELIMINATING_VULNERABILITIES_IN_GLOBAL_FULFILLMENT_NETWORKS/links/68d1e8cb11d348252ba6db60/ZERO-TRUST-SECURITY-AND-MFA-DEPLOYMENT-AT-SCALE-ELIMINATING-VULNERABILITIES-IN-GLOBAL-FULFILLMENT-NETWORKS.pdf
18. Dennstädt, F., et al. (2024). *Implementing large language models in healthcare while managing risks*. *NPJ Digital Medicine* (or Nature partner journal), 2024.
19. Bajwa, S., & authors. (2021). *Artificial intelligence in healthcare: transforming the ...* (PMC review). *BMC/PMC* 2021.
20. Adari, V. K. (2024). The Path to Seamless Healthcare Data Exchange: Analysis of Two Leading Interoperability Initiatives. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(6), 11472-11480.
21. Archana, R., & Anand, L. (2025). Residual u-net with Self-Attention based deep convolutional adaptive capsule network for liver cancer segmentation and classification. Biomedical Signal Processing and Control, 105, 107665.
22. Bussu, V. R. R. Leveraging AI with Databricks and Azure Data Lake Storage. https://pdfs.semanticscholar.org/cef5/9d7415eb5be2bcb1602b81c6c1acbd7e5cdf.pdf
23. Raju, L. H. V., & Sugumar, R. (2025, June). Improving jaccard and dice during cancerous skin segmentation with UNet approach compared to SegNet. In AIP Conference Proceedings (Vol. 3267, No. 1, p. 020271). AIP Publishing LLC.