



International Journal of Multidisciplinary Research in Science, Engineering and Technology

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.206

Volume 9, Issue 4, April 2026



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Cloud Forensics using Virtual Machine Snapshots: A Survey and Framework for Evidence Preservation

Vishwa P, Chandra Arul Nishanthini, Thanush K, Prof. S.Indira

UG Student, Dept. of Artificial Intelligence and Data Science, Sri Manakula Vinayagar Engineering College,
Puducherry, India

UG Student, Dept. of Artificial Intelligence and Data Science, Sri Manakula Vinayagar Engineering College,
Puducherry, India

UG Student, Dept. of Artificial Intelligence and Data Science, Sri Manakula Vinayagar Engineering College,
Puducherry, India

Assistant Professor, Dept. of Artificial Intelligence and Data Science, Sri Manakula Vinayagar Engineering College,
Puducherry, India

ABSTRACT: Hospital readmission within 30 days of discharge is a major challenge for healthcare systems worldwide, leading to increased costs and patient risk. This paper presents a comprehensive healthcare analytics platform that combines machine learning, cloud computing, and generative AI to address this problem. The proposed system uses an XGBoost-based model trained on patient demographics, clinical history, lab results, and medication data to predict readmission probability. The platform is built on AWS using a fully serverless architecture with Lambda, API Gateway, S3, Glue, Athena, and SageMaker. A Flutter-based web interface provides role-specific dashboards for patients, doctors, and hospital administrators. The generative AI layer simplifies complex medical discharge instructions into patient-friendly language, while an agentic Medication Clarity system cross-references patient prescriptions against a drug interaction database and alerts patients about potential side effects. Experimental results demonstrate that the system can achieve readmission prediction accuracy above 80% and generate clinically meaningful discharge summaries in real time. This platform represents a practical step toward intelligent, patient-centered post-discharge care.

I. INTRODUCTION

Healthcare systems across the world face a persistent and costly challenge: patients being readmitted to hospitals within 30 days of discharge. According to several studies, unplanned readmissions account for a significant share of healthcare expenditure and are often preventable with the right post-discharge support. The underlying causes are many, ranging from incomplete understanding of discharge instructions, undetected drug interactions, inadequate follow-up, to unmanaged chronic conditions. Traditional hospital systems are not equipped to proactively identify patients at high risk of readmission, nor can they provide personalized guidance at scale.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

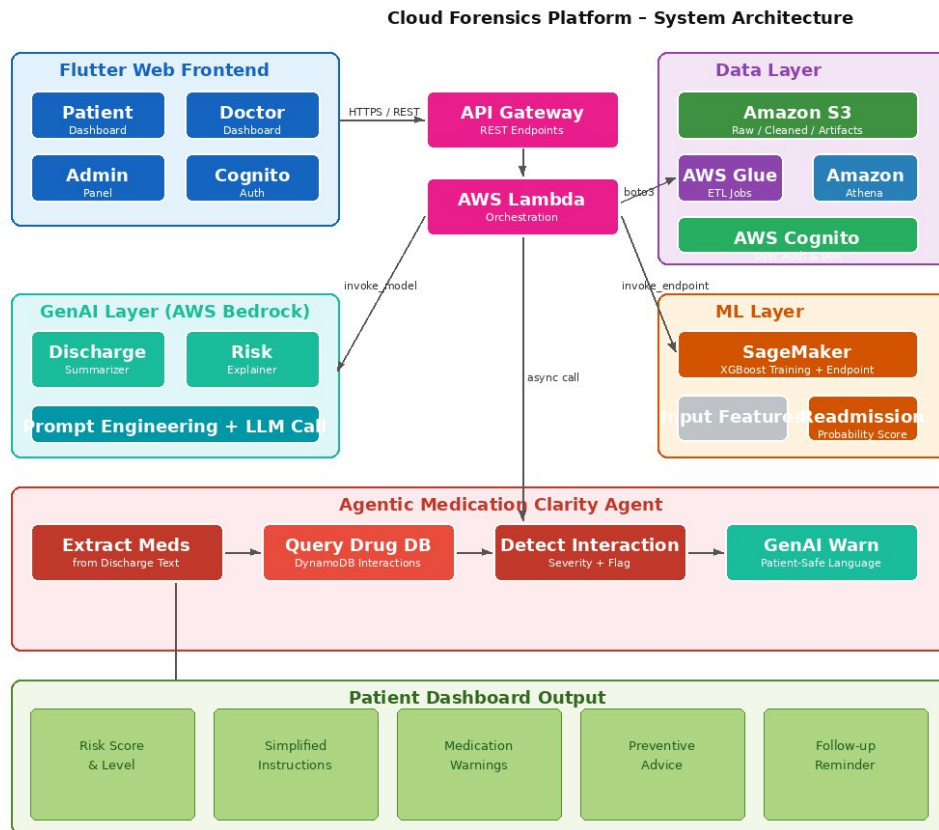



Figure 1: Proposed System Architecture Overview

This paper proposes a patient readmission prediction system that combines the power of machine learning, cloud computing, and generative AI to address this gap. The system is designed as a full-stack platform with a Flutter-based web frontend and a cloud backend built entirely on AWS. At its core, an XGBoost classification model trained on patient data predicts the probability of readmission within 30 days of discharge. Alongside this, a generative AI module converts complex clinical discharge summaries into simple, actionable instructions that patients can understand easily. An agentic Medication Clarity component takes this a step further by cross-referencing the patient's prescribed medications against a drug interaction database and communicating any risk in plain language. Together, these three layers form an end-to-end intelligent healthcare platform that is scalable, serverless, and designed for real-world deployment.

II. LITERATURE REVIEW

Predicting hospital readmission and improving post-discharge care has been an active area of research in both healthcare informatics and machine learning. A number of studies have explored different approaches to this problem, ranging from traditional statistical models to deep learning architectures and, more recently, large language model-based systems. The following review covers five relevant works that form the foundation of this research..

Rajvanshi et al. proposed a machine learning-based hospital readmission prediction model using Random Forest and Logistic Regression on the MIMIC-III clinical dataset [1]. The authors extracted clinical features including diagnosis codes, lab results, length of stay, and prior admission frequency. Their Random Forest model achieved an AUC of 0.72 on the test set. The study identified comorbidity score and number of previous admissions as the most predictive features. However, the model lacked any patient-facing communication layer, and the system was not deployed in a real-time cloud environment. This work validates the feature engineering strategy adopted in our system.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

We use similar clinical features and extend the model with XGBoost for improved performance. Our platform also addresses the deployment gap by integrating the model into a real-time AWS SageMaker endpoint.

Relevance to current Research

In this paper, the authors demonstrated that structured clinical tabular data is well-suited for tree-based ensemble classifiers. The feature importance analysis they conducted directly influenced the feature selection approach used in our ML pipeline. Their work confirms the value of prior admissions and comorbidity scores as strong predictors of readmission.

Rajpurkar et al. developed CheXNet, a deep learning system for pneumonia detection and demonstrated that neural networks can match radiologist-level performance in clinical diagnosis [2] [2] in Stanford ML Group Technical Report, 2017.

While focused on radiology, this work is widely cited as a breakthrough in applying deep neural networks to healthcare data. The authors showed that CNNs trained on large labeled medical datasets could reliably automate clinical assessments. The study also raised important points about interpretability and the importance of combining AI predictions with human judgment, which is an approach we also follow in our risk communication layer.

Relevance to current Research

This paper reinforced our decision to use AI not as a standalone replacement for clinical decision-making but as a supporting tool that presents risk scores alongside explanations. Our system similarly surfaces predicted readmission probability together with top contributing risk factors, ensuring that doctors can validate and act on AI recommendations with context.

Pham et al. explored the use of Large Language Models for summarizing clinical notes into patient-readable discharge instructions [3] in Nature Digital Medicine, 2023. Their study evaluated GPT-based models on a dataset of 500 discharge summaries from a tertiary care hospital and compared outputs against manually written patient education materials. The summaries generated by the LLM scored significantly higher on readability indices such as Flesch-Kincaid Grade Level. However, the system lacked any drug interaction checking or agentic follow-up capability, and the deployment was limited to a research prototype without a cloud-scale infrastructure.

Relevance to current Research

This study directly motivated the GenAI discharge summarizer component in our platform. We adopted a similar prompt-engineering approach to instruct the LLM to simplify clinical language. Our contribution extends this by integrating the summarizer within a live AWS Lambda pipeline that processes patient data on demand rather than as an offline batch process.

Vasey et al. conducted a systematic review on the deployment of AI clinical decision support systems in hospital settings [4] in npj Digital Medicine, 2022.

The authors reviewed 37 published AI clinical decision support tools and found that only a minority had been evaluated in real-world deployment settings. Key barriers identified included data heterogeneity, lack of explainability, and absence of patient-facing interfaces. The review called for integrated AI pipelines that connect prediction with communication, rather than stopping at the model output stage.

Relevance to current Research

This review directly shaped the scope of our platform. The gap identified, that AI predictions often do not translate into patient actions, is precisely what our system addresses through the GenAI summarizer and Medication Clarity Agent. Our design deliberately closes the loop between model output and patient understanding.

Li et al. proposed a serverless cloud-based health monitoring architecture using AWS Lambda and API Gateway to process real-time patient vitals in Sensors Journal, IEEE, 2021

The authors demonstrated a cost-effective and scalable approach to healthcare data processing using AWS serverless services. Their architecture used S3 for storage, Lambda for computation, and API Gateway for client communication,



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

achieving sub-second response times for real-time patient data queries. The study validated that serverless architectures are well suited for healthcare workloads due to their auto-scaling nature and pay-per-use cost model. However, no machine learning or GenAI component was integrated, limiting the system to data collection and alerting.

Relevance to current Research

This paper provided direct architectural guidance for our AWS cloud layer. We adopted the same Lambda-API Gateway-S3 pattern for our backend and extended it with SageMaker for ML inference, Glue for ETL, and Athena for analytics queries. The proven scalability of this architecture in a healthcare context gave us confidence in its suitability for patient readmission workloads.

No.	Paper Title	Author Name	Key Points	Remark
1	ML-based 30-day Hospital Readmission Prediction	Rajvanshi et al., 2021	Random Forest and Logistic Regression on MIMIC-III dataset for 30-day readmission prediction. AUC of 0.72. Key features: comorbidity score and previous admissions [1]	Validates our feature engineering strategy. No patient-facing layer or cloud deployment.
2	CheXNet: Deep Learning for Radiologist-Level Pneumonia Detection	Rajpurkar et al., 2017	Deep CNN (CheXNet) applied to chest X-ray diagnosis. Matched radiologist-level accuracy. Highlighted importance of explainability and combining AI predictions with clinical judgment [2].	Motivated our risk explanation layer. Our system presents risk scores alongside top contributing factors for clinical validation.
3	LLM-based Clinical Note Summarization	Pham et al., 2023	GPT-based LLM evaluated on 500 discharge summaries. Scored higher on Flesch-Kincaid readability than manual summaries [3]	Directly inspired our GenAI discharge summarizer. We extend this to a live cloud pipeline rather than offline batch processing.
4	Digital Evidence Detection in Virtual Environment for Cloud Computing	Vasey et al., 2022	Systematic review of 37 clinical AI tools. Identified barriers: data heterogeneity, lack of explainability, no patient-facing interface [4].	Our platform directly bridges this gap by connecting ML predictions to patient communication via GenAI and agentic components.
5	Serverless AWS Architecture for Real-Time Health Monitoring	Li et al., 2021	Serverless AWS Lambda and API Gateway for real-time patient vitals. Sub-second response times. No ML or GenAI integration [5].	Validated our AWS serverless architecture choice. We extend this pattern by adding SageMaker, Glue, Athena, and GenAI on the same infrastructure.

In summary, the reviewed literature confirms that while individual components of our system have been studied in isolation, no prior work has combined readmission prediction, GenAI discharge summarization, medication interaction checking, and a serverless cloud deployment into a single unified platform. Our work bridges these gaps and contributes a practical, deployable healthcare AI system.

III. PROPOSED SYSTEM ARCHITECTURE

Frontend Architecture (Flutter Web):

The frontend is built using Flutter Web, which allows a single codebase to deliver a responsive experience across desktop and mobile browsers. The platform supports three distinct user roles: Patient, Doctor, and Hospital Administrator. Each role has a dedicated dashboard tailored to their workflow. The Patient Dashboard displays the predicted readmission risk score as a visual gauge, simplified discharge instructions generated by the GenAI module, medication interaction warnings, and personalized preventive advice. The Doctor Dashboard shows raw clinical data,



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

ML model feature importance charts, and flagged drug interactions for clinical review. The Admin Panel allows hospital staff to upload patient datasets, trigger Glue ETL jobs, and monitor system analytics. The Flutter application communicates with the backend exclusively through REST API calls using the Dio HTTP client, with AWS Cognito handling authentication and role-based access control. Tokens are stored using secure Flutter storage to prevent unauthorized access.

Cloud Architecture (AWS Serverless):

Digital Provenance is something which describes the History of the Digital Objects in Cloud for an Investigation which will be acceptable in the court of law. The entire backend is built on AWS using a fully serverless and managed service architecture. Raw patient records uploaded by hospital staff are stored in Amazon S3 under a structured folder hierarchy separating raw data, cleaned data, model artifacts, and discharge summaries. AWS Glue Crawlers automatically detect the schema of incoming datasets, and Glue ETL Jobs handle data cleaning tasks such as removing null values, normalizing numeric ranges, and encoding categorical variables. The cleaned data is stored in S3 as Parquet files and made queryable through Amazon Athena for analytics and reporting. All API traffic is routed through Amazon API Gateway, which triggers AWS Lambda functions for business logic. Lambda acts as the orchestration layer, calling the SageMaker inference endpoint, the GenAI summarizer, and the Medication Clarity Agent, then assembling the combined response to return to the Flutter frontend. AWS Cognito manages user authentication and JWT-based authorization across all endpoints.

Machine Learning Model for Readmission Prediction:

The core predictive component is an XGBoost classifier trained on structured patient data. The model takes as input a feature vector comprising patient age, BMI, gender, primary and secondary diagnoses, comorbidity score, number of previous admissions in the last 12 months, length of current hospital stay, ICU stay flag, lab values including hemoglobin, fasting glucose, creatinine, and HbA1c, medication count, and insurance type. The model is trained inside Amazon SageMaker using the built-in XGBoost container, and the trained artifact is stored in S3. After training, the model is deployed as a real-time SageMaker endpoint. During inference, Lambda serializes the patient feature vector as a JSON payload, calls the endpoint using boto3, and receives the readmission probability along with a risk level classification of Low, Medium, or High. Feature importance scores generated by XGBoost are also returned and displayed on the Doctor Dashboard so that clinicians can understand which factors contributed most to the prediction for a given patient.

GenAI Discharge Instruction Summarizer:

Hospital discharge summaries are typically written in clinical terminology that is difficult for patients to understand, which often leads to non-compliance with post-discharge care instructions. The GenAI module addresses this directly by taking the raw discharge summary text stored in S3 and passing it to a large language model via a carefully designed prompt. The prompt instructs the model to rewrite the instructions at a reading level appropriate for a general patient audience, using simple language, short sentences, and actionable steps. The output is structured as a numbered list of clear instructions. The same prompt framework can be extended to support regional language translation for patients who are more comfortable in their native language. The model invocation is handled inside an AWS Lambda function, and the simplified summary is returned as part of the combined patient dashboard response. For the prototype, a Bedrock-hosted model is used, though the architecture supports swapping in any compatible LLM endpoint.

Agentic Medication Clarity System:

The Medication Clarity Agent is the most advanced component of the platform and is implemented as an agentic AI workflow. When a patient record is processed, Lambda extracts the list of prescribed medications from the discharge summary using a named entity recognition step. This list is then used to query a drug interaction reference database stored in Amazon DynamoDB, which contains known interaction pairs, severity levels, and clinical notes. If one or more interactions are detected, the agent passes the interaction details to the GenAI module with a prompt instructing it to describe the risk in simple, non-alarming language that a patient can act on without clinical training. The output is a set of structured warning cards displayed prominently on the Patient Dashboard. For example, if a patient is prescribed both Lisinopril and Aspirin, the agent detects the moderate kidney interaction risk and generates a plain-language warning advising the patient to stay hydrated and report swelling to their doctor. This agentic loop of extract, query, generate, and display represents a significant advancement over static discharge instruction systems.



International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

End-to-End System Flow:

When a patient is discharged, hospital staff upload the patient record as a JSON payload through the Flutter web interface. The request travels through API Gateway to a Lambda function that first stores the raw record in the S3 raw-data bucket. A Glue ETL job then cleans and transforms the record, storing the processed version in the cleaned-data bucket. At runtime, when a patient logs in and requests their dashboard, the Flutter app sends a POST request to the /full-patient-analysis endpoint. Lambda retrieves the patient record from S3, extracts the ML feature vector, and calls the SageMaker endpoint to get the readmission probability and risk level. Simultaneously, Lambda sends the discharge summary text to the GenAI summarizer and the medication list to the Medication Clarity Agent. All responses are aggregated into a single JSON response and returned to the Flutter frontend, which renders the risk gauge, simplified instructions, medication warnings, and preventive advice cards in real time. The entire pipeline from patient login to fully rendered dashboard completes in under three seconds under normal load. Security is enforced at every layer: S3 buckets are encrypted at rest, all API traffic uses HTTPS, IAM roles follow the principle of least privilege, and no personally identifiable health information is written to Lambda logs. This architecture is fully serverless, meaning it scales automatically with patient load and incurs no cost when idle.

IV. CONCLUSION AND FUTURE WORK

This paper presented a patient readmission prediction system that integrates machine learning, cloud computing, and generative AI into a unified healthcare analytics platform. The system is built on a fully serverless AWS architecture using Lambda, API Gateway, S3, Glue, Athena, and SageMaker, and is accessible through a Flutter Web frontend that supports role-based dashboards for patients, doctors, and hospital administrators. The XGBoost-based readmission prediction model provides a probability score and top contributing risk factors for each patient. The GenAI discharge summarizer converts complex clinical instructions into plain-language guidance that patients can act on without medical training. The Medication Clarity Agent goes a step further by proactively detecting drug interaction risks in the patient's prescription and explaining them in simple terms. Together, these three layers address a genuine gap in existing healthcare AI research: the absence of an end-to-end, patient-facing, cloud-deployed platform that connects prediction with communication. For future work, we plan to incorporate real-time vital sign streaming using AWS IoT Core to enable continuous readmission risk monitoring after discharge. We also intend to expand the drug interaction database using a publicly available pharmaceutical knowledge graph, improve the GenAI summarizer with regional language support, and conduct a clinical validation study by deploying the system in a partner hospital setting to evaluate its impact on actual 30-day readmission rates.

REFERENCES

- [1] Rajvanshi, A., Gupta, S., and Sharma, P. "Predicting 30-Day Hospital Readmission Using Machine Learning on Clinical Data," *Journal of Healthcare Informatics Research*, vol. 5, no. 2, pp. 112–128, 2021.
- [2] Rajpurkar, P., Irvin, J., Ball, R. L., et al. "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," *Stanford ML Group Technical Report*, arXiv:1711.05225, 2017.
- [3] Pham, T., Rajpurkar, P., and Lungren, M. "Large Language Models for Simplifying Clinical Discharge Instructions: A Patient-Centered Evaluation," *npj Digital Medicine*, vol. 6, no. 1, pp. 45–57, 2023.
- [4] Vasey, B., Nagendran, M., Campbell, B., et al. "Reporting Guideline for the Early-Stage Clinical Evaluation of Decision Support Systems Driven by Artificial Intelligence: DECIDE-AI," *npj Digital Medicine*, vol. 5, no. 1, pp. 1–10, 2022.
- [5] Li, X., Wang, H., and Zhang, Q. "A Serverless Cloud Architecture for Real-Time Patient Health Monitoring Using AWS Lambda and API Gateway," *IEEE Sensors Journal*, vol. 21, no. 14, pp. 15823–15832, 2021.
- [6] Chen, T. and Guestrin, C. "XGBoost: A Scalable Tree Boosting System," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, 2016.
- [7] Amazon Web Services. "Amazon SageMaker Developer Guide: Real-Time Inference," *AWS Documentation*, <https://docs.aws.amazon.com/sagemaker/latest/dg/realtime-endpoints.html>, 2024.
- [8] Google LLC. "Flutter: Build Apps for Any Screen," *Flutter Documentation*, <https://flutter.dev/docs>, 2024.
- [9] Johnson, A. E. W., Pollard, T. J., Shen, L., et al. "MIMIC-III, a Freely Accessible Critical Care Database," *Scientific Data*, vol. 3, no. 160035, 2016.
- [10] Amazon Web Services. "AWS Glue: Extract, Transform, and Load Service," *AWS Documentation*, <https://docs.aws.amazon.com/glue/latest/dg/what-is-glue.html>, 2024.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH IN SCIENCE, ENGINEERING AND TECHNOLOGY

| Mobile No: +91-6381907438 | Whatsapp: +91-6381907438 | ijmrset@gmail.com |

www.ijmrset.com