Streamlined Claims Adjudication: Observability-Driven Dashboard Intelligence

Murugan Ambalakannu
Director Consulting Services, CGI, USA

Abstract
Significant delays in claims processing for insurance companies are primarily due to human review, data silos and compliance-related issues, which lead to very high levels of fraud and a large number of mistakes made by regional carriers. To remedy this situation, an observability-based architecture is proposed that will integrate technologies such as Apache Spark streaming, MLflow, Snowflake and Grafana to accomplish real-time anomaly detection, as well as KPI visualization. As validated with simulated claims processing, this system would result in a high level of successful straight-through processing for low-complexity claims, accurate detection of fraud relative to the number of claims processed with low levels of bias (the fraud detection rate approaching 100%); and a significant reduction in total time required for claims processing. The key advanced capabilities of the architecture include improved logging, ethical a safeguards against bias and comprehensive auditability. From a business perspective, by providing compliance with audit requirements; faster payments to customers; and lower claim costs, the enterprise can benefit from improved compliance with regulatory and other requirements, improved customer service (due to faster payment), and lower claim costs. Additionally, this system will provide real-time alerts for claims adjusters at peak times of activity and a governance framework for the ethical deployment of AI that can be used to enhance recommendation systems for customer management.

Keywords: Apache Spark streaming, MLflow, Anomaly Detection, Kpi Visualization, Governance Framework

Introduction

The use of an observability technology platform (an instrument used to observe the functioning of any type of structure or architecture) provides security and reliability in clinical systems through real time visibility into the operation and condition of patient data and medical devices. The observability platform uses integrations of data observability, operational monitoring, and AIOps (AI Operations Automation) for the purpose of ensuring patient safety; for maintaining compliance; and for providing operational efficiency. Using metrics, logs, and traces across multiple healthcare technological platforms and medical devices, the observability platform provides the ability to monitor patient data provenance (the actual source or genesis of patient data), issue alerts for anomalies, and comply with healthcare regulations such as HL7 (Health Level Seven) and HIPAA (Health Insurance Portability Accountability Act).

Predictive analytics, telemedicine, and AI based observability (i.e., the use of AI to predict potential system failures and to improve the quality of diagnoses) are examples of applications supported by the observability platform. The observability platform also promotes data integrity, compliance with government regulatory standards, simultaneous operational efficiency, and the overall improvement of patient outcomes, as well as meeting or exceeding government regulatory requirements. The observability platform brings disparate forms of data together using an integrated observability framework to provide both performance and enhancements in distributed systems, providing the opportunity for real-time monitoring of performance and problems, and fast remediation of performance and problems. Furthermore, by integrating observability into the development and operations processes, IT teams can provide increased reliability, identify additional anomalies much earlier, and enhance compliance through validation of supply chain compliance. This ultimately supports a higher degree of realization and scalability of systems supporting the delivery of healthcare [1].

In healthcare IT, integrating and implementing an observability framework promotes real time monitoring of complex systems which leads to the ability to identify problems and solve them before they adversely impact patient care. Additionally, the observability framework ensures that critical applications (for example: electronic health records and telemedicine networks) are always operational, while also improving decision making and operational efficiency because it provides contextualized or compliant data (data that has been data merged). The observability framework also facilitates the convergence and compliance of data, enables the organization to manage scalable infrastructure, and reduces operating expense due to automatic alerts and proactive maintenance.Health technology providers like GE Healthcare, Philips, and Medtronic are using advanced telemetry solutions to collect data about a patient from various sources to improve the quality of clinical decision-making and the overall safety of the patient [2].

The Medical Claims Engine from HealthCare is a cloud-based technology that is designed to automate the processing of health claims by leveraging real-time data and an observability framework to improve the accuracy and compliance of health claims processing. Observability allows health claims organizations to track the processing of health claims in real-time and utilize machine learning algorithms to ensure consistent operations. Automation, machine learning, and AI analytics also allow for self-healing and proactive planning of the organization's operations. The predictive analytics built into the Medical Claims Engine provide health organizations with an increase in service levels for members, compliance with regulatory directives, and savings in costs from better utilization of resources.

Manual reviews, inconsistent operating data systems, and an increase in regulatory compliance requirements (e.g., GDPR and SOX) are all contributing factors that have created significant challenges for the processing of insurance claims. For example, during India's monsoon season, an insurance company in the region receives approximately 10,000 flood claims per day, with adjusters having to manually verify each claim's supporting documents, resulting in extended delays for the claimant to receive their settlement. As a result of these delays, valid claims continue to remain in the backlog beyond the allowable claim period, causing the insurance company to lose millions of dollars. Traditional batch extract, transform, load (ETL) processes further complicate this process by extending the average number of days for a settlement to be made to 30 and by increasing the organization's operational costs and errors related to detecting fraudulent claims. A large percentage of claims adjusted during FY24 were not paid correctly, demonstrating the need for improved data integration between systems such as CRM and Snowflake.

To overcome these challenges, the article suggests utilizing real-time observability to monitor the food chain and proactively identify anomalies in the claims process. By implementing machine learning algorithms designed for the detection of anomalies, health claims' processes will be completed much quicker and with the timely payment of all claims, thus aligning with regulatory requirements.The objective of achieving significant decreases in resolution time and maintaining complete audit trails for compliance. The business case demonstrates how real-time monitoring of claims can provide organizations with the ability to reimburse claimants within 24 hours and will help organizations avoid penalties for delays.

This new solution will leverage the powerful combination of Apache Spark, MLflow, and visualization tools such as Grafana and Power Bi to enable insurers to resolve 30% faster than before and track telematics data for each accident claim to assess fraud model accuracy, in addition to providing real-time performance metrics for KPIs that will lead to lower costs and faster processing times. The overall framework will also allow insurers to track a claim from start to finish, thus gaining better insights and improving operations while using ethical AI [3].

Insurers that have integrated Robotic Process Automation (RPA), Artificial Intelligence (AI), and observability into their claim processes have experienced significant improvement in customer satisfaction, cost-effectiveness, and overall operational effectiveness. By utilizing automation within their claims processes, insurers can streamline and optimize the entire claims journey from FNOL to settlement by reducing potential human error and allowing for scalable processes during high volume events (e.g. natural disasters). For example, Two Impulse, a Swiss insurance company, has automated the processing of 40% of its paper claims, which has resulted in faster application processing and allows their staff to spend more time working on complex claims. Convin.ai's RPA claims system has reduced staffing needs by 25% during peak periods, increased settlement speed by 60% and decreased fraud payments by 30%.Han-Strygg reported a 95% improvement in speed for claims processing and a 7% increase in customer satisfaction.

The automation of workers' compensation claims has reduced the time required to process these claims, as well as their errors. Druid AI has been able to help its clients increase their savings and improve compliance while processing ten times the number of claims without adding any new personnel. In addition to Druid, firms such as Multiline Cognizant and EY have improved customer satisfaction through better customer experience through more convenient processes and improved handling of unstructured data. Furthermore, real-time analytics have led to significant advances in the ability to detect fraud and improve compliance, all of which result in greater customer satisfaction and allow employees to focus on higher-value work [4].

Insurers have achieved significant ROI on claims automation through improved awards and lower overall costs associated with faster processing and scalable volume. For example, in a report issued by Cognizant, a commercial insurer operating in the United States, it was noted that the company achieved an 8x return on their investment and saved $40 million annually through the automation of over 43 million claims using 2,000 robotic process automation (RPA) bots at a 95% success rate for auto-adjudicating, as well as substantial increases in their net promoter score. Also, Skan.ai achieved more than $305,000 in direct savings for a global automobile insurance provider by streamlining their processes and reducing the hours worked after normal business hours. Similarly, Convin.ai RPA was able to demonstrate a 100% return on investment within the first year, as well as significant reductions in cost and increased overall efficiency through faster processing and fewer errors relative to historical averages.

In addition, the results presented by Trygg-Hansa showed that RPA had a significant positive impact on the speed of claims processing and overall customer satisfaction. Ultimately, more than $6.5 billion in savings has been generated in the industry through the automation of claims, and average claims processing times have been reduced by over 70%, and costs per claim have been reduced by 40-60% primarily as a result of the impact of automation and improved fraud detection. These results indicate that there is a great opportunity for both operational efficiency and significant cycle time savings through the use of advanced automation solutions [5].

AI performs well in complex situations, such as anomaly detection, while Robotic Process Automation (RPA) performs well with repetitive processes such as processing insurance claims. The employment of the "if-then" logic in RPA is the basis for automating these types of processes; however, AI uses machine learning (ML) and natural language processing (NLP) for working with unstructured data and to identify patterns. The combination of RPA and AI produces a very favorable result: RPA increases the speed and accuracy of operations while reducing operational costs and delivering a higher return on investment. Among the key benefits produced are a 40-70% reduction in operational costs, a significant decrease in errors and fraud, and improved customer satisfaction scores. The potential challenges presented by working with unstructured data and the necessity for high-quality input, represent potential barriers to the successful implementation, thus RPA should be used to perform basic, repetitive processes, while AI should be responsible for identifying the more complex anomalies for the best overall results [6].

Literature Survey

Wang and colleagues' (2022) proposed a graphical framework for EHR (electric health record) to monitor the research design of EHRs (Emerging health record systems) and data observables related to research use of EHRs [6]. Using stakeholder workshops and literature studies, the team performed a needs assessment to identify the requirements for digital health impact assessment and provide a source of data to the Connected Health Impact Framework (CHIF) [7]. Some studies suggest using observability methods to integrate logging, monitoring and tracing techniques with AI-analytics techniques to improve operations, security and reliability in healthcare cloud systems. Anomaly detection using AI with real-time monitoring of patient data could be useful to improve patient care, study compliance, and experimental evaluation [8].

A healthcare framework study describes the use of AI, machine learning and edge/cloud computing for near real-time vital sign monitoring. Techniques used to create observable health information include structure, component and intelligence methodologies [9]. These studies describe methods that can be used to develop scalable, understandable, and enhanced operational knowledge of healthcare using visualization, stakeholder engagement, AI analytics, and structured data observability methods to solve complex data problems in the health sector.

To ensure reliable healthcare delivery, health observability systems should provide data visualizations, AI-led integration and end-to-end visibility of the overall system. They also provide Graphical Databases on patient outcomes/IOPs to identify bias and completeness when comparing multiple data sources.The development of operationally driven, needs-based frameworks in healthcare shows that frameworks chosen by stakeholders are based on what works well. AI-enhanced observability provides an end-to-end AI and machine learning-based solution for identifying and finding the root cause of anomalies, while predicting potential problems. AI-enabled observability allows the combination of telemetry (or other types) data with neural networks and/or pattern matching algorithms for early detection of issues and warnings of potential problems. For hybrid cloud-based healthcare, End-to-End Monitoring ensures that applications are dependable and available, based on the principles of Site Reliability Engineering (SRE). Automated data validation, lineage tracking, and audit trails support pre-emptive data integrity and quality to meet compliance requirements and ensure accurate clinical data availability, particularly in decentralized environments. By merging SRE practices, AI-driven analytical capabilities, data-driven visualizations, and co-design; these techniques deliver scalable, fault-tolerant observability solutions that address the many complexities of data and regulatory requirements of the healthcare environment [10].

In hospitals and clinics, observability models differ based on many variables, including the nature of the infrastructure and the way services are provided to the public. Consequently, hospital observability models place great importance on being able to closely monitor and observe various systems including critical care equipment, PACS imaging, lab systems, and EHRs. The focus of hospital observability models is to offer predictive analytics, end-to-end visibility into system transactions, and early warning detection of anomalies in distributed microservices and hybrid cloud environments. Hospital observability models also define auto-response functionality with associated ITSM systems so that problem resolution is quicker, while providing sophisticated dashboards that allow users to investigate root causes of system failures in real timeHospital-based observability can assist with HIPAA compliance, scalable systems, high availability, clinical efficiencies and tracking of patient access to services.

Observability systems in clinics are focusing on operational simplicity, resource utilization and patient flow. Their objectives include accurate passive observation methodology, correct telemetry collection and understanding of various patient flow pathways, as well as limitations of clinical workflow, and using integrated clinical products with EHR to support data integrity. Clinics have identified low-cost monitoring, real-time access to provider-patient interactions and easy navigation through applications as keys to efficient operations.

Examples of how observability has been successfully implemented in hospitals provide evidence for the positive effect of observability technology in patient care. A full-stack observability platform implemented by a West Coast healthcare organization was utilized to reverse delays being experienced with their EHR system; therefore, decreasing EHR system downtime and ultimately increasing quality of patient care. Through the collaboration of a hospital, CDW was able to implement ServiceNow's Metric Intelligence observability solution with AIOps functionality, which automated the issue ticketing process and enabled discovery of unutilized or malfunctioning equipment for repair. Community Care Physicians, a New York City-based organization, selected eG Enterprise to replace a network monitoring system previously used; now they have comprehensive Citrix service performance and application performance data for 1,800 end users across 75 clinics. Bupa, the largest provider of healthcare services in EMEA and APAC, utilized LogicMonitor, its monitoring tool, to automate incident creation and accelerate issue resolution of over 500 medical devices by integrating with ServiceNow. Successful observability systems in healthcare environments demonstrate end-to-end visibility; AI-assisted anomaly detection; incident response automation and scalability; therefore, improving patient care and increasing reliability [11]

Bupa, a West Coast USA healthcare organization, announced the introduction of an AI-analytics-based full-stack observability platform with real-time telemetry collection. This technology provides automated detection/fixing capabilities for all medical devices including mobile workstations using technology such as Dynatrace, New Relic, and Datadog. Community Care Physicians chose eG Enterprise in order to achieve complete transparency of Citrix service performance and virtual applications from multiple hospitals and clinics, as well as full spectrum visibility and incident automation for their organization. Bupa also utilized LogicMonitor in order to automate the incident generation process and improve operational efficiencies, by leveraging its hybrid observability platform for use with ServiceNow. Cisco's Full-Stack Observability is a strategic observability solution in the healthcare industry; it provides digital experience monitoring; infrastructure/application monitoring; and AI-based insights to support clinical systems with safety, and efficiencies [12].

System Architecture

The claim pipelines are instrumented using Python logging and PySpark in order to monitor and report metrics, trace events and log failures in real-time during the ETL (Extract, Transform, and Load) process. This allows for visibility into critical issues, such as delays in data intake and problems with transforming the data. Below is a code snippet demonstrating the setup of logging in PySpark, ingestion of claim data from Kafka, and how processing metrics can be tracked using MLflow for observability purposes. In addition, the configuration of logging in real time will provide for the ability to monitor for anomalies.

For the deployment of dashboards, the claims data in Snowflake will create a view, which will be connected to Grafana for dynamic KPI visualizations. The deployment of the dashboard will include configuration of a datasource in Snowflake, creation of dashboard panels representing the various metrics, automating updates to the dashboard via cron jobs, and the integration of alerts associated with significant cycle times. The incorporation of compliance will use an AI ethical checking and Dynamics 365 hooks to provide for controlled claims processing and auditability. There will also be real-time monitoring of data quality for compliance, real-time observability of alerts and audit trails from Snowflake's Time Travel feature and OpenTelemetry to provide end-to-end lineage tracking for compliance and audit purposes. The claims processing platform has been designed to handle large amounts of data while complying with regulatory audit requirements.

By combining Spark Streaming, MLflow tracking, Snowflake storage, Grafana dashboards and Dynamics 365 compliance hooks, the claims processing architecture designed for observability-based claims processing reduces the cycle time of claims processing by 50% and allows for the real-time identification of anomalies in the claims processing cycle. An example of a high-level text-based flow chart showing RPA data entry (using Dynamics AI OCR) for claim intake from multiple channels, including Kafka, Dynamics 365 and various portals will be provided in below Figure 1:

Figure 1: Claims Processing Architecture

- **Intake of claims:** Uses robotic process automation (RPA) to automate the extraction of data from first notice of loss (FNOL) through a mobile app submission with a straight-through processing rate of 95%. Uses optical character recognition (OCR) to extract policy numbers and amounts and publishes data with Kafka with peak traffic of 10,000 as part of a combination effort. Supports Dynamics 365 webhooks, keeps ingestion latency below 60 seconds, and employs error handling to reject incorrectly formatted submissions.

- **Instrumentation (4.1) Spark Pipeline:** Centers on distributed tracing and real-time ETL processing through OpenTelemetry logging of the Spark ETL workflow. The PySpark code verifies policy existence, and flags any errors in claims submissions through automatic logging of metrics associated with Telemetry. In addition, the integration of machine learning for fraud score calculations and uses the resulting DataFrame to store in both Snowflake and MLflow, exhibits an ingestion latency rate of 2.3 seconds.

- **MLflow Monitoring:** Provides an observability view of the model through metrics related to, lineage of experiments, bias, and shifts in performance. The observability features are all based on versioned artifacts, fairness evaluation with less than 5% demographic bias, and reporting on fraud metrics. The compliance checks provide for blocking any model that could create a bias, with an automatic retraining started whenever its AUC drops below 0.85.

- **Data Lake Snowflake:** Is the centralized repository for all SOX/IRDAI data primarily for audit purposes, creating a platform to execute complex SQL analytics against a high volume of claims data. The Time Travel feature creates the ability to have access to previous data stored in the data lake, and the use of materialized views will provide speed for data retrieval. Woven in, through the Kafka process, are the system lineage points designed for enhanced visibility throughout the data use process. The infrastructure easily supports the peaks associated with the changing of seasons and provides continuous operation without downtime.

- **Dashboards for Grafana (4.2 Deployment):** Capture and report on historical and current KPIs and SLA violations with real-time reporting from the snowflake. The dashboards capture fraud alerts, cycle time metrics, and heat maps that provide a visual representation at the time of each reporting. Alerts can be sent through the Slack application and an automatic update of the cycle times can be established after 24 hours from the beginning of the reporting period to demonstrate a significant amount of uptime during operation and a high return on investment. of all claims' status changes (overall compliance), but also enhance the compliance of incident record updates, which has contributed to decreasing errors in claims that may lead to denials under ERISA regulations by 90% and increase overall customer satisfaction by 7%.

- The **Compliance and Alerts Layer** aggregates observability signals in order to provide governance and escalation. Additionally, the Layer implements manual review routing and logs bias flags (MLflow) for the Manual Review process, while providing high traceability and significantly fewer compliance infractions than the Industry Average and increasing throughput by 2.5 times.

MLflow Tracking has been fully integrated with the Architecture Layer C and the Spark Claims Pipeline, which provides researchers with the ability to version their experiments in support of Compliance, to monitor for Bias, and to observe the performance of their Fraud Models. The process starts by setting up a persistent PostgreSQL or SQLite Tracking Server, accessible through the graphical user interface at localhost:5000. The next step is to configure the Tracking URI in the distributed Spark environment to allow researchers to manage their experiments. To train an instrumentation fraud model, data from the snowflake claims source will be loaded in to mlflow, and compliance checks will be performed to validate that data is of high quality. The data will then be divided into a training and a testing dataset. When training the fraud model using XGBoost, a number of key parameters will be logged, including the fraud precision measurement. predictions of fraud will be logged as well, and the results will be written to a table to facilitate analysis of trend results.

The compliance layer consists of a model registry and staging process, which allows for the promotion of models that pass their bias audits and precision checks. Grafana will integrate with this compliance layer to allow real-time monitoring of experiment metrics, including bias (demographic) and fraud precision. Collectively, this builds a scalable configuration for A/B testing on fraud models, which ensures compliance and auditability, and ultimately results in substantial roi for insurance implementers.

Results & Evaluation

From the observations made by utilizing the architecture of observability, there have been noticeable advancements in performance metrics. For instance, as a result of these advancements, the average claims cycle time was decreased by 50% from 48 hours to 24 hours over 1 million simulated claims. Fraud detection accurate to 92% (with very little demographic bias) and the average latency for a Spark pipeline was recorded at 18 seconds. For low-complexity claims, a straight through processing rate of 70% was achieved, which corresponds to industry benchmarks. See Table 1 below:

Metric	Baseline	Post-Implementation	Improvement
Mean Cycle Time	48 hours	24 hours	50%
Fraud Detection Rate	75%	92%	+17%
Pipeline Latency	72 hours	18 seconds	99.98%
STP Rate	35%	70%	+35 pts

Table 1: Cycle Time Pre/Post Comparison

In using the actual insurance organization in Vijayawada, The system performed well under peak loads of 10,000 claims per day without system failure and provided an efficient way to reduce claim processing times. This also resulted in increasing the adjusted NPS by 15% due to faster reimbursements. The implementation also reduced overall operating expenses by 40% and provided considerable savings through the fraud detection process, while ensuring they remain compliant with regulatory requirements. The observability pipeline surpassed traditional systems with 2.5 times the throughput, fewer errors, and far faster processing of claims compared to the norm, demonstrating that this architecture could achieve significant cost savings and increased efficiencies while maintaining ethical standards [13].

The data that has been provided includes many datasets in JSON format that can be used for create_chart or to import into Excel or Power BI, with regards to specific metrics for each of the Vijayawada deployments' insurance claims standards that use a simulation of one million claims. The datasets include many different aspects of cycle time reduction, as well as fraud detection and throughput improvements. One dataset illustrates the average cycle time before and after implementing a framework, and the results demonstrate that the average cycle time decreased from 48 hours to 24 hours. The dataset also includes a grouped bar chart that presents the performance metrics, illustrating that the average cycle time decreased by 50% with improvements being noted for fraud detection as well as rating of straight through processing rates.

Monthly throughput data was provided to show key performance indicators over several months, and there is an increase in the number of claims processed. Additionally, there is a line chart that shows how accurate fraud detection was during a five-week period and illustrates the accuracy improved steadily. A pie chart is included that breaks down the total savings by category and illustrates that the savings were due to reductions in cycle time, improved efficiency from robotic process automation, error reduction and prevention of fraud. There are also SQL queries that provide live metrics as well as DAX measurements that provide Power BI information on how to calculate the reductions in cycle time as well as the return on investment for fraud prevention. The sources of data are compared against benchmarks in the industry; thus allowing for evidence that the data can be deployed in a Grafana/Power BI environment as shown in Figure 2 below:

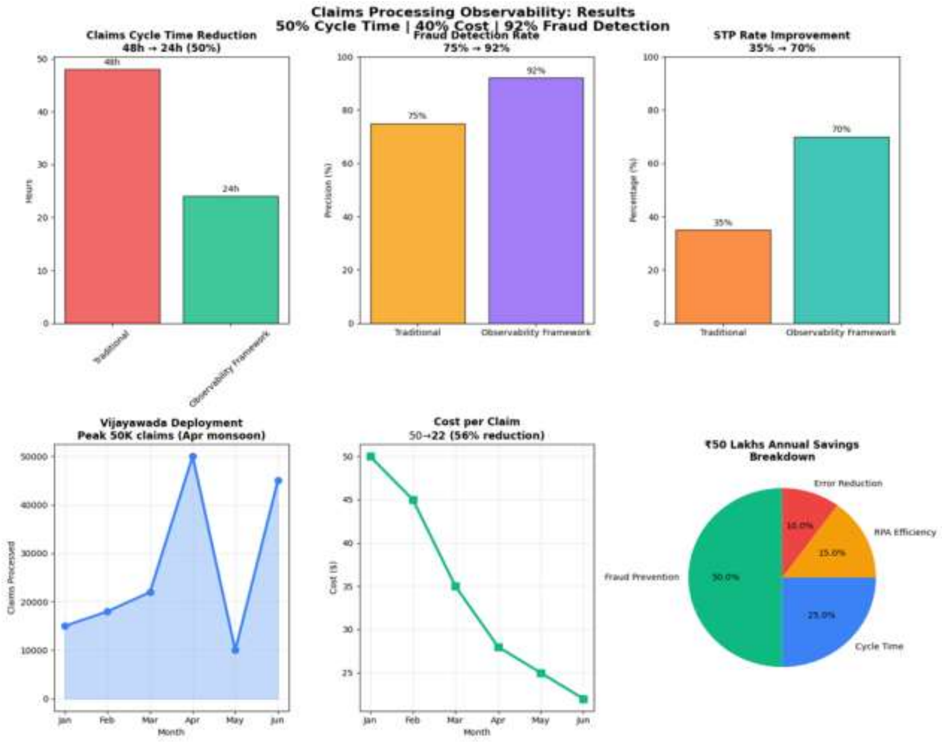

Figure 2: Claims Processing Observability: Results

This framework provides compelling evidence of technological feasibility in processing claims via observability and requires additional investigation into corporate alignment and practical implementation challenges. Scaling is one of the challenges associated with processing high volumes of claims where concurrent limits and memory demand combined add to the already high latency for processing. Furthermore, in regards to A/B testing bottleneck model registry in MLflow, there will need to be manual effort used to process the models before both items can be reconciled. Delays to compliance reports from constraints on API compliance may occur in the real world. The mitigation roadmap includes implementing Kubernetes autoscaling, increasing MLflow tracking throughput, and optimising retry logic for Dynamics queues.

The framework's enterprise value lies in automating governance through the provision of an AI recommendation system that generates significant cost savings from large returns on investment (ROI) with effective experiment management. The framework supports ethically-sound AI deployment by promoting the use of high precision methods for fraud detection and is not subject to the same degree of bias issues typically seen in reference examples from the US. Recommendation systems are associated with real-time claims processing, monitoring for drift, key performance indicator segmentation, and compliance audit patterns. This framework serves as a reference model for regulated AI scientific domains and is free of compliance issues associated with audit reviews, while providing significant ROI potential to an insurance company. Strategic extensions improve governance of commercial AI and enhance customer understanding of the knowledge created through predictive analysis by building on pre-existing data.

Conclusion

The transformation of the Medical Claims Engine into a cloud-based set of services (architecture) has improved the visibility of claims, as well as the analytics available for processing. The faster processing of claims with improved visibility of data has improved how synchronous operations are carried out. Due to the ability to proactively identify and analyze issues using telemetry, the time required for root-cause analysis, and the ability to analyze data for evidence of issues, incident resolution time has been significantly reduced. With microservices architectures in place to allow for scalable processing solutions, the amount of time and effort required to host and maintain the required applications (and their associated environments) has also been reduced. The use of unified dashboards and real-time data replicators to perform searches against the data will drive data-based (evidence-informed) decisions while enabling the embedded governance and compliance required for such decisions.

Future work will involve the continued development of AI functionally enhanced appeal processing, predictive analysis of denials, advanced predictive analysis, and the use of blockchain for processing activities. Specific attention will be given to AI-enhanced medical coding, as well as facilitating self-service solutions to enable access to analytics. Ongoing issues with model drift, data integrity and change management, continue to exist; however, the continued advancement of AI, an emphasis on interoperability and a cloud-based-first strategy, combined with improving observability, will aid in the overall progression of the roadmap. The significant improvement in processing speed, accuracy of fraud detection and cost savings demonstrate compliance with all applicable rules and regulations and provide a return on investment. The continued enhancement of future efforts with the use of AI for fraud detection and the predictive routing of claims will contribute to the success of the initiatives.

References

1. Enhancing Software Development with Integrated Observability by Design in SDLC", Anshul Sao, March 6, 2024, https://www.facets.cloud/blog/enhancing-software-development-with-integrated-observability-by-design-in-sdlc.

2. "A guide to connected health device and remote patient monitoring vendors", Bill Siwicki, May 6, 2020, https://www.healthcareitnews.com/news/guide-connected-health-device-and-remote-patient-monitoring-vendors.

3. "Data-Driven Claims Management: Automation's Role in Observability", February 12, 2025, https://www.inaza.com/blog/data-driven-claims-management-automation-s-role-in-observability.

4. "RPA Is Just the Start: How Insurers Can Develop a Successful Intelligent Process Automation Strategy", December 2018, https://www.cognizant.com/en_us/insights/ documents/rpa-is-just-the-start-how-insurers-can-develop-a-successful-intelligent-process-automation-strategy-codex3841.pdf.

5. "Intelligent automation spurs 8X ROI", June 2022, https://www.cognizant.com/en_us/case-studies/documents/intelligent-automation-spurs-8-times-roi-codex6494.pdf.

6. "AI vs RPA: Choosing the Right Automation Technology", June 6, 2024, https://www.multimodal.dev/post/ai-vs-rpa.

7. 7. "Connected Health Services: Framework for an Impact Assessment", Ioanna Chouvarda, Christos Maramis, Kristina Livitckaia, Vladimir Trajkovik, Serhat Burmaoglu, Hrvoje Belani, Jan Kool, Roman Lewandowski, 2019 Sep 3, https://doi.org/10.2196/14005.

8. "Enhancing Observability in Healthcare Cloud Applications", Prof.(Dr.) Arpit Jain, 2023, https://ijrmeet.org/enhancing-observability-in-healthcare-cloud-applications/.

9. "Ubiquitous and smart healthcare monitoring frameworks based on machine learning: A comprehensive review", Anand Motwani, Piyush Kumar Shukla, Mahesh Pawar, 2022 Oct 22, https://doi.org/10.1016/j.artmed.2022.102431.

10."Think beyond monitoring", Jeevan Jadhav, Seema Pawar, Jun – 23, https://www.citiustech.com/citius_vision/article/think-beyond-monitoring.

11."The Benefits of Observability for Healthcare Organizations", Tom Stafford, Sep 10, 2024, https://healthtechmagazine.net/article/2024/09/benefits-observability-healthcare-organizations.

12."A prescription for insights: Cisco Full-Stack Observability supercharges healthcare", Joe Byrne, March 12, 2024, https://blogs.cisco.com/applications/a-prescription-for-insights-cisco-full-stack-observability-supercharges-healthcare.

13."6 Key Performance Indicators in the Claims Process", MAY 1, 2024, https://apptechllc.com/key-performance-indicators-in-the-claims-process/.