The Structural Shift: How Unified Claims Platforms Reshape Payer Economics

Nagabhushanam Bheemisetty

Abstract
Healthcare claims processing is a significant operational cost for payers, yet most systems used are old and do not meet today's high-volume, multi-regulatory environment. In this article, we propose a related strategic framework to transition healthcare organizations from fragmented, manual claims operations to unified, intelligent claims platforms. We provide evidence from recent implementations as well as the academic literature to make the case that consolidating legacy data systems, using machine learning to detect/fix fraud and adjudicate claims, and automating workflow orchestration can dramatically change how much it costs to do administrative work. The transformation proposed — built upon five foundational capabilities — will yield 15–30% lower administrative cost, 30–40% lower denial rework volumes, and positive ROI within 18–30 months. In addition, this transformation will transform claims processing into a strategic asset rather than a reactive liability, allowing for scalable growth, regulatory robustness, and an improved experience for key stakeholders.
Keywords: Claims Processing Automation, Healthcare Administrative Costs, Machine Learning Fraud Detection, Intelligent Adjudication, Unified Data Architecture, Value-Based Care Enablement

1. Introduction

A conflict exists between two aspects of the healthcare industry in the U.S. today; increasing the complexity of medical practice and regulations while the foundation of how a claim is processed was developed decades ago (when medical care was not near as complicated), thus leading to inefficiencies in the payment of claims. Medical claims processing is a major financial component of running a health plan or payer; whether it be a 'stand-alone' payer, self-insured employer or government payer, each has a significant cost associated with processing every claim and additional cost to an already burdened system of reworking denied claims [1].

The complexity and the additional layering or fragmentation of claims data into various software systems or legacy adjuication engines, third party administrator(s), provider systems and external clearinghouses causes multiple non-value added activities to take place, therefore increasing the operating expenses of the respective health plans or payers while degrading provider and member satisfaction. Because of the manual nature of claims adjudication, the use of static rule engines and inconsistent policy interpretation creates loops of rework that incur substantial costs throughout the claims lifecycle. Clearly, there is a substantial financial burden associated with this inefficiency. In fact, many estimates place the average administrative cost per claim for all payers (commercial and governmental) at $2.00 to $4.00 and an average denial reimbursement/administrator employee ratio of 30% [5].

Additionally, fraud, waste and abuse adds another 3% to 5% as a percentage of total claims dollars (used to pay for services) that are realized through retrospective detecting of the fraud, waste and abuse. At the same time, prior authorization processes continue to be reactive in nature, generating documentation and delaying the decision. As medical care is being inflated at a higher pace than health plans/payers can adjust to, the inability to process claims efficiently, consistently and at a large volume are directly increasing the medical loss ratio; which will continue to erode the operating margins of the payer. The conclusion reached from the assessment of this problem is that the solution is not going to be to further optimize existing processes but to radically redesign the architecture of the claims process by moving from fragmented, workflow driven processes to consolidated, real-time claims intelligence platforms [3][8].

2. Methodology

The goal of this article is to provide a comprehensive framework based on empirical research for effectuating an evolution in the industry. The project consists of four different, but interconnected, aspects. First, we review the literature available on digitization of claims by looking at both benefits/costs associated with electronic claims systems in isolated cases and compiling a global database of these data points through academic research [1][3].

As part of the ongoing process of completing the literature review above and obtaining available examples of successful applications of machine learning techniques for fraud detection, we have completed a literature search for peer-reviewed articles on machine-learning based fraud detection techniques including regression models and anomaly detection models. We collected papers on data preparation, training, and testing of regression and anomaly detection techniques against actual claims data [2].

To illustrate the successful application of AI to claims processing and financial handling, we found case studies of:

1) Government insurance agencies receiving an ROI of 9x through automation of their invoice processing,
2) Health Plans reducing the processing time for appeals by 30% through the use of AI and robotic process automation and
3) Third party administrators receiving over $13 million in total savings through execution of their digital transformation projects.

Additionally, we assessed modern public and private healthcare organizations that have implemented process automation and/or AI as part of their claims/financial handling processes [4][5].

Finally, we used contemporary approaches to develop a technical architecture for a data platform to incorporate medallion architecture design approaches that are applicable for processing healthcare transactions. The financial model presented includes baseline assumptions for mid-tier national payer organizations (e.g., 50M claims per year) and applying assumptions for administrative costs per claim ($2-$4) as well as the percent of expense to the administrative cost (12-15%). These assumptions, along with performance factors that have been previously published, were applied to create a financial model of a payer organization.

3. Workflow Transformation: From Fragmented to Unified Operations

3.1 The Structural Inefficiency Problem

According to the analysis, disconnected claims architectures create systemic cost multipliers resulting in indirect processing costs. Core systems that are siloed force multiple validations for each transaction from various data sources, thus creating discrepancies between different valid transactions, requiring manual intervention for reconciling. The use of manual adjudication processes, static rule engines and inconsistent interpretations of company policies leads to the creation of denial rework loops whereby 10% - 15% of total claims must be rewritten and/or appealed and/or reviewed a second time, thereby artificially inflating the administrative expense ratio and placing strain on relationships between providers and payers [5].

A prime example of this problem is prior authorization processes. Document-intensive and reactive in nature, this delays timely decisions for patient care resulting in friction in the provider satisfaction survey score, as well as limiting the member's retention of all patients. In addition, fraud detection is primarily conducted after funds were expended and as such, recovery processes are completed after the payment has already been made rather than proactively preventing fraud in the adjudication processes. Thus, the claims processing system represents a cost multiplying-system facing increasing scrutiny by regulators [8].

The financial implications of operational failures only serve to exacerbate the challenges of operations. Each additional day spent adjudicating a transaction impacts predictability of cash flow. Each percentage point increase in avoidable rework due to denials increases administrative overhead. Each instance of inconsistent application of rules increases the likelihood of non-compliance. An analysis further illustrates that digital transformations for the claims adjudication process can reduce cycle time for processing claims and improve speed to reimbursements while decreasing the rate of denials; however, this must be viewed as an investment in the entire health care system and not as discrete technology projects.

3.2 Evidence from Implementations

Implementation examples show substantial measurable results can be produced via targeted transformation. A national injury insurance program utilized an AI-driven method of processing invoices for which the investment returned nine times what was spent on implementing this method, and man-hours invested in completing this task were decreased by as much as 90%. The amount of time required to process an invoice dropped from many days to less than 24 hours, resulting in 100% accuracy of the data that was extracted, which allowed workers to devote their time to tasks of higher value as opposed to data entry roles [2].

Within the appeal arena, a Medicare Advantage program serving over 10 million members implemented AI and robotic processing automation to reduce high costs, low accuracy levels, and long duration times incurred as a result of processing appeals. By processing appeals faster as a result of this intervention, the duration for processing appeals fell by approximately 30%, leading to approximately 4% higher Star Ratings (contributing to financial performance) and a 75% decrease in the total number of complaints received by Medicare holders due to improved accuracy when processing appeals [10].

Fraud detection applications also present a similar opportunity. Utilizing machine learning algorithms that have been fed historical claims information, the algorithms are capable of identifying patterns associated with potential fraudulent activity (behaviors) with a greater level of accuracy than traditional, rule-based methods. The integration of additional data sources in conjunction with the ability to frequently adapt models will migrate fraud detection from a retrospective recovery-based function to a real-time prevention-based function [3][8].

A third-party administrator recently changed from a manual-based operation to a digital-based operation and utilized a number of different solutions across the three areas of claims editing, billing review, and payment processing. The transformation resulted in over 300,000 transactions that migrated from a paper to a digital format, streamlined workflows, and saved in excess of $13 million [5].

3.3 Unified Platform Architecture

To redesign the architecture so that it includes all of the components of the claims process within one system (a unified claim intelligent platform) instead of fragmented legacy systems will solve many of the common systemic inefficiencies. Figure 1 below shows the target state architecture, which consolidates and centralizes the entire set of claims-relevant data into one single fabric layer (a unified fabric): therefore allowing for real-time policy governance, intelligent adjudication of claims, predictive fraud detection, and conditional workflow orchestration [9].

International Journal of Scientific Research in Engineering and Management (IJSREM
Volume: 10 Issue: 03 | March - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Figure 1: Target-State Unified Claims Intelligence Platform Architecture

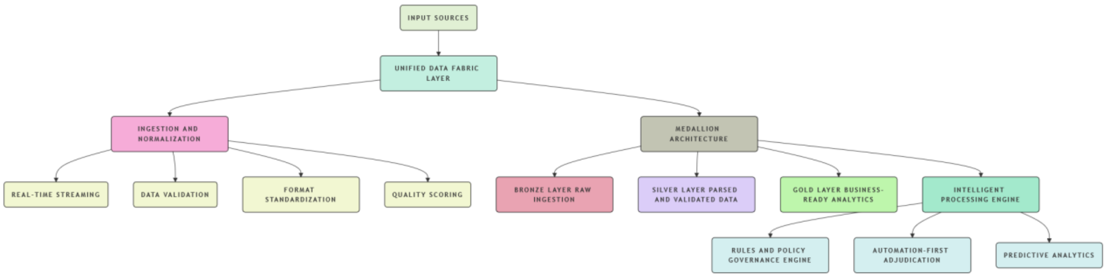

The architecture consists of five layers that interact and work together:

- **Input Sources Layer:** Input sources include many data streams going into the platform: claims submitted from provider systems (837), eligibility data from member web portals, pharmacy data (prescription fills), policies/benefits admin rules, and claims sent through external clearinghouses between entities (routing).
- **Unified Data Fabric Layer:** this is a new foundational shift from having fragmented data to having a consolidated data source (intelligence). It normalizes data as it enters (real time ingestion) and provides data validation by structuring/standardizing the formats for all sources of data and scoring data quality. The data stored within this Architecture (medallion architecture) have 3 layers/stages of information [9]:

  - ✓ **Bronze Layer** – includes raw data ingested (retaining original format from all sources) to maintain auditability.
  - ✓ **Silver Layer** – includes transformed, validated, and consolidated data (providing consistent member, provider, and claims view).
  - ✓ **Gold Layer** – includes aggregated data that are business-ready (for analytic/reporting purposes)

- **The Intelligent Processing Engine:** This engine consists of three engines that execute the core claims intelligence:

  - ✓ **Rules and Policies Governance Engine -** this engine contains a centralized repository for rules to manage, real-time updates, rules applied by jurisdiction, and a complete audit trail.
  - ✓ **Automation-First Adjudication Engine -** this engine contains the Decision Intelligence (claim) models for determining how claims should be processed (i.e., straight-through processing of low-risk claims, automated prior auths for evidence-based cases, or human-reviewed queue for complex scenarios).
  - ✓ **Predictive Analytics Engine -** this engine contains the Machine Learning models for detecting potential fraudulent patterns, identifying anomalies, analyzing provider risk profiles over time, and flagging high cost claims prior to payment.

- **The Workflow Orchestration Layer:** the claims received are assigned an appropriate risk profile through a Conditional Execution Engine. The Conditional Execution Engine routes the claim for processing, based on the risk profile, and ensures compliance with all policies/rules related to the claim. This layer eliminates unnecessary manual steps and provides a controlled governance process.

- **The Output Channels:** It provides remittance advice (835 transaction) to the provider for processed claims, provides the member with an explanation of benefits, payments that were executed, provides full audit trails for all claims processed, and analytic reporting.

The Integration of the multiple layers creates a claims architecture that is free from redundancy of manual actions and can grow to scale compliance and automation-first claims operations. Current implementations utilize the framework for processing multiple types of transactions including, but not limited to, 837 Healthcare claims, 835 Payment Advices, 278 Preauthorization Requests, and Eligibility Transactions; all of these transactions must also maintain a secure file exchange using encryption with all trading partners.

3.4 Quantified Performance Metrics

The document details metrics to measure the improvement in the three areas of finance, operations, and customer experience after transition. The metric for finance indicates the ability to lower administrative cost per claim (from $2 - $4 to $6.40 - $9.60), as well as reduce the leakage of fraud, waste, and abuse from 3 - 5% to 2.25 - 3.75%. The medical loss ratio is estimated to improve 1 - 2 percentage points. The operational metrics indicate that denial/rework rate is expected to decline from 10 - 15% to 6 - 9% while at the same time there has been a large decrease in average cycle time of claims (from 10 - 18 days to 6.5 - 12 days). The customer experience metrics demonstrate an expectation for an increase in provider reimbursement times as the current baseline NPS for providers is expected to increase 10 - 15 points and provide for a more accurate estimate of member out-of-pocket costs. See Figure 2 below:

Figure 2: Multi-Dimensional Performance Comparision

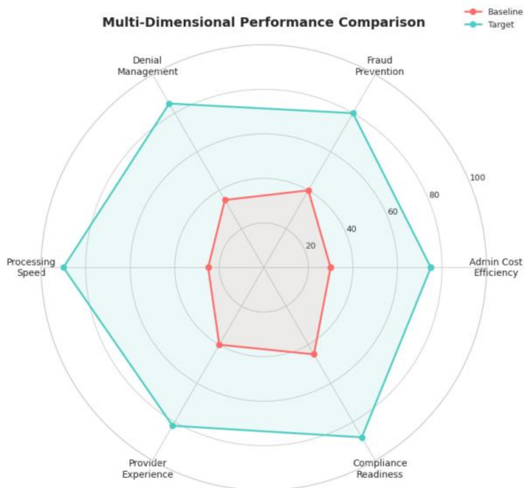

Metrics show a complete review of improvement targets and baseline positions after transformation across three categories; Finance, Operations, and Experience. Financial metrics include a decrease in cost per claim from $2 to $4 down to $6.40 to $9.60 (20% decrease) and a medical loss ratio improvement of 1-2%. Operational metrics include an average claim denial/rework rate drop from 10-15% to 6%-9% and a dramatic reduction in average claim cycle time from 10-18 days to 6.5-12 days. In addition, it is estimated that appeal volume will decline by 30%. Experience metrics illustrate an acceleration of provider reimbursement 6 to12 days and an anticipated increase of 10 to 15 points on the Net Promoter Scores (NPS) for providers, and improvements in the accuracy of member out-of-pocket expense estimates can be seen in Table 1 below.

Table 1: Baseline vs. Projected Performance Metrics (3-Year Horizon)

Metric Category	**Baseline (Current State)**	**Target (Post-Transformation)**	**Improvement**
Financial Metrics			
Administrative Cost per Claim	\$2–\$4	\$6.40–\$9.60	20% reduction
Administrative Expense Ratio	12–15% of premium	9–12% of premium	15–30% improvement
Fraud, Waste & Abuse Leakage	3–5% of claims spend	2.25–3.75% of claims spend	25% reduction
Medical Loss Ratio	Baseline	1–2 point improvement	1–2 percentage points
Operational Metrics			
Denial/Rework Rate	10–15% of claims	6–9% of claims	30–40% reduction
Straight-Through Processing Rate	20–30% (estimated)	50–80%	30–50% increase
Average Claims Cycle Time	10–18 days	6.5–12 days	25–35% reduction
Appeal Volumes	Baseline	~30% decrease	Based on implementation evidence
Experience Metrics			
Provider Reimbursement Speed	10–18 days	6–12 days	Accelerated
Provider NPS	Baseline	+10–15 points	Projected
Member Out-of-Pocket Estimate Accuracy	Variable	Improved	Predictive

4. Discussion

4.1 Foundational Capabilities for Transformation

Research indicates that there are five key foundation characteristics needed for successful claims transformation:

- **Unified Data Fabric:** Consolidate all the claims, provider, member, pharmacy, authorizations and policy data into an integrated, interoperable data layer that provides a unified data experience. This will provide real-time ingestion and normalization of the data, and resolve all the reconciliation gaps. The use of the Medallion Architecture pattern has proven out as an effective implementation template utilizing the three layers of bronze for raw ingestion, silver for parsing and validating the data, and gold for performing business analytics resulting in one source of truth across all adjudication processes .
- **Intelligent Rules & Policy Governance:** The move from static rules configurations to dynamic rules configurations in a centralized architecture that supports real-time policy changes and jurisdictional application of the rules, automated audit trails and transparent version control of the rules has ensured consistency in the business rules used for adjudication and instant alignment with regulatory requirements

without re-deploying a new system. This addresses inconsistent policy interpretation which drives claims denial and increases compliance exposure [6].

- **Automation First Adjudication:** Utilizing Decision Intelligence models will create straight-through processing of low-risk claims, predictive fraud detection at submission, automated prior authorization approvals for evidence based claims and risk-based routing of claims for human review only when necessary. Evidence shows that more than 70% of the total claims processed in high volumes and with low complexity can be successfully processed without human interaction when implemented properly [7].
- **Predictive Risk & Fraud Analytics:** The incorporation of Machine Learning models that will allow for anomaly detection prior to payment, identification of claim trajectories that have a high likelihood of becoming high-cost claims and identification of risk patterns that will emerge in the provider community will transition fraud detection from a retrospective recovery basis to a proactive prevention basis. Gradient Boosted Models such as CatBoost and Anomaly Detection techniques such as Isolation Forest have shown to produce excellent results on real-life claims data.
- **Workflow Orchestration & Transparency:** Implementing end-to-end workflow automation that is integrated with Decision Engines to facilitate conditional execution based on Risk Profile, Compliance Requirements and Policy Rules will remove non-value added manual steps but will allow for governance control. By eliminating manual touch-points, this will significantly reduce administrative expenses for all payer participants.

4.2 Implementation Considerations

Transformation's success is impacted by many factors. First, target high-frequency, low-complexity claims to generate the highest return on investment (ROI) in the shortest time frame by pursuing current volumes that use maximum time by adjusters with the highest potential for automation . Second, speed to integrate is beneficial-proven automation can be deployed within months, not years if solutions are already integrated with core systems rather than replacing them.

Third, a continuous feedback loop from models is critical for increasing automation rates over time. Machine learning models must also evolve to respond to fraud scheme changes and operationally to shifts in workflow. Fourth, provider-type and context's varied benefit-cost ratio indicate optimization must be based on readiness for infrastructure, scale of operation, and management capability/ownership.

Cost implication of transformation also needs to be considered. Significant initial commitments are needed for software development, hardware, infrastructure and training staff; ongoing costs will also be required, which include cyber-security, systems maintenance and user-support. However, benefit-cost analysis indicates there is a strong economic basis to support transformation; one national study approximated benefit-cost ratio of 90.06 (incremental) for implementation of collaborating provider within United States health system.

4.3 Strategic Implications

Unified claims platforms not only provide operational & financial metrics, but they also provide strategic advantages too. Unified claims platforms can provide a scalable ability to grow without having to have equal administrative growth, which is crucial with healthcare growing in volume and margins getting less. They provide a unique way to differentiate from competitors based on speed and transparency as provider and member expectations continue to grow. Unified claims platforms can reduce exposure to compliance risk by implementing automated audit trails and ensuring consistent rule application in the claims process.

Most importantly, a modernized claims infrastructure provides a platform for value-based care models and innovations related to bundling payments. When claims data is clean, timely, and accessible, it allows organizations to move from just transactional processing to managing population health; measuring performance against benchmarks; reducing overall claim expense while improving patient care, as illustrated below.

Figure 2: Strategic Impact Pyramid

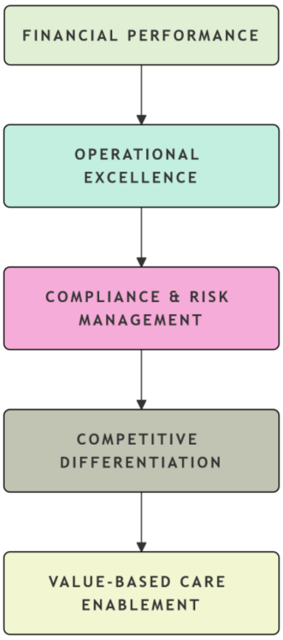

5. Limitations and Future Research

The limitations of this synthesis represent some challenges that will have an impact on its findings. In this respect, the digital transformation literature contains only examples from a limited number of organizations that had experience in using digital processes early on. While financial projections can give some insight into possible returns for an organization, the impact of digital transformation may be very different depending on the circumstances of each organization as well as the marketplace in which they operate.

Further research is needed to examine the value to organization of digitizing administrative functions, to do cross-national comparisons of the effectiveness of digital transformation strategies across countries, and to combine time motion and financial analyses to assess the financial impact of existing process inefficiencies. As electronic claims systems are put into place in Ghana, Rwanda, and Tanzania, evaluation of these systems is extremely important in terms of their ability to reduce delays in processing payments, to enhance the reliability of payment processes, and to foster trust in the providers of health services.

6.Conclusion

Current healthcare claims processing is structured in a manner that creates a structural liability. The current operating model is designed for transactional insurance environments, which are now struggling under the sheer volume of claims, the tremendous complexity of today's claims (compliance, for example), and growing regulatory pressure. The increasing amount and quality of evidence available suggests a new path forward: from fragmented, workflow-driven claims operations to unified, intelligent claims platforms based on a) consolidated data fabrics; b) machine learning-enabled adjudication; and c) automation-first workflows.

Measuring the results of early implementations show promise: 30-40% reduction in denial rework; 25% improvement in fraud prevention; 30-50% higher rates of straight-through processing; and positive ROI within two years. Collectively, these improvements will lead to improved margins; decreased risk of regulatory issues; decreased time for providers to receive reimbursement for services; and laying the groundwork for value-based reimbursement models.

The transition from traditional claims processing to modernized claims processing will require organizations to view claims processing as a strategic capability rather than simply a back-office function. To those organizations that successfully make this transition, they will place themselves in a position to experience resilient growth in an increasingly difficult and challenging healthcare environment.

7. Acknowledgement

This analysis was made possible due to the contribution of all the healthcare payers, technology partners, and academic scholars who participated in the implementation of these organizations, and whose experience and empirical research was used to inform this analysis. We want to specifically acknowledge and thank the organizations who have supplied transformation results in an open manner, which provides evidence-based guidelines for the entire industry.

References

1. da-Costa Vroom FB, Shimada Y, “Digitizing Health Insurance Claims: Building Sustainable Pathways to UHC”, Johns Hopkins Bloomberg School of Public Health, Center for Global Digital Health Innovation. 2025 Aug, https://publichealth.jhu.edu/center-for-global-digital-health-innovation/digitizing-health-insurance-claims-building-sustainable-pathways-to-uhc.
2. Sean Mitchell, “AI cuts invoice processing time at Queensland injury agency”, IT Brief Australia. 2025 Jun, https://itbrief.com.au/story/ai-cuts-invoice-processing-time-at-queensland-injury-agency.
3. P Ashok, Abhijit Sambhaji, “DurgeFraud Detection and Prevention in Healthcare Insurance Claims Using Machine Learning Regression Models”, IEEE Xplore. 2025 Jun, https://doi.org/10.1109/ICDSBS63635.2025.11031751.
4. Azure Fabric X12 EDI Processing Pipeline, “Healthcare transaction processing with medallion architecture”, GitHub. 2025 Sep, https://github.com/vincemic/ai-fabric-etl.
5. “ACS & Zelis: A smarter way to save and scale”, Zelis Case Studies, 2025 Mar, https://www.zelis.com/case-studies/client-focused-partnership-drives-innovation-and-savings/.
6. “Optum launches AI system to speed medical claims”, Emily Olsen, Healthcare Dive, Oct. 22, 2025, https://www.healthcaredive.com/news/optum-real-ai-speed-claims-review-united-health/803448/.
7. “Top 5 DACH Insurer Automates 70% of P&C Claims end-to-end”, omni:us Case Studies, 2025 Sep, https://omnius.com/case-studies/no-touch-pc-claim-automation/.

8. Chengamma Chitteti; Mopuri Yamuna; Mattam Srinath; Chakali Govardhan; Alavalapati Vignatha, "Healthcare Insurance Fraud Detection Using Machine Learning", 30 May 2025, https://doi.org/10.1109/ICOEI65986.2025.11013497.
9. "Overview of CMS claims data transformations in healthcare data solutions", Microsoft Learn, 2025 Mar, https://learn.microsoft.com/en-us/industry/healthcare/healthcare-data-solutions/claims-data-transformations-overview?toc=%2Findustry%2Fhealthcare%2Ftoc.json&bc=%2Findustry%2Fbreadcrumb%2Ftoc.json.
10. "AI and RPA Transform Appeals, Reducing Cost by 30%", Sagility Health Case Studies, 2025 Oct, https://sagilityhealth.com/case-studies/ai-and-rpa-transform-appeals-reducing-cost-by-30/.