

# Denial Automation in Healthcare Apps

A step-by-step guide for engineering and product teams building denial automation into revenue cycle workflows

Published by [Code-B Solutions](#) | 2026 | Healthcare Technology Series

## Foreword

The blog that accompanies this whitepaper makes the case for why denial automation matters.

This document does not repeat that case. If your team has already aligned on the problem of rising denial rates, payer AI asymmetry, recoverable revenue sitting uncaptured, what you need next is a concrete plan for building the capability.

***This playbook is written for the engineers, product managers, and technical leads who will actually implement it.***

It covers the six implementation phases in sequence, the failure modes at each stage, and the decisions that determine whether a denial automation system compounds value over time or stalls after the first deployment.

## Phase 1: Data Readiness Audit

*Before writing a line of automation logic, you need to understand what your data actually looks like.*

Most implementation timelines are blown not by engineering complexity but by data quality problems that nobody assessed upfront. This phase should take two to three weeks and is non-negotiable.

### What to audit:

Pull your last 24 months of 837P/837I (claim submission) and 835 (Electronic Remittance Advice) transaction files from your clearinghouse or EHR data warehouse.

You are looking for five things:

**First**, payer ID consistency. The same insurer frequently appears under dozens of different payer IDs across legacy systems.

Before any ML model can learn payer-specific denial patterns, those IDs need to be normalized to a canonical identifier. A single BCBS plan appearing as 14 different payer codes is not unusual. Map them all.

**Second**, CARC/RARC code completeness. Claim Adjustment Reason Codes and Remittance Advice Remark Codes are the structured vocabulary your denial categorization engine will operate on.

If your 835 data has incomplete or missing remark codes on a significant portion of records common in older EHR configurations you have a gap that will degrade downstream automation. Quantify the percentage before proceeding.

**Third**, linkage between submissions and remittances. Can you reliably join your 837 submission records to their corresponding 835 ERA records?

If your clearinghouse and EHR store these in separate systems with inconsistent claim identifiers, the join logic needs to be built before any feature engineering begins.

**Fourth**, clinical data availability. For the NLP appeal layer to function, clinical notes and medical necessity documentation need to be retrievable from the EHR by date of service and claim ID.

Confirm that your FHIR API surface (or HL7 v2 interfaces) can return DocumentReference and DiagnosticReport resources for historical encounters. If you are on Epic, test your Bulk FHIR export against a sample date range now.

**Fifth**, COVID-era distortion. Claims data from 2020 and 2021 has denial pattern anomalies caused by emergency billing rule changes, telehealth expansion, and payer policy suspensions.

If those years are in your 24-month window, flag them and consider excluding them from model training while keeping them in validation analysis.

**Deliverable from this phase:** A data readiness scorecard that quantifies record completeness, linkage integrity, and the estimated engineering effort to clean and normalize the dataset. This becomes the honest baseline against which your implementation timeline is set.

## **Phase 2: Denial Taxonomy & Prioritization**

**Not all denials are worth automating first. Choosing the wrong starting point is the most common reason denial automation programs lose momentum.**

Before configuring a single bot or training a single model, your team needs to produce a denial taxonomy specific to your organization and rank it by automation ROI.

**Build the taxonomy in two dimensions:**

- Volume tells you where the queue is clogged. Dollar value tells you where the revenue is.

- These are rarely the same list. A high volume of low-dollar administrative denials (missing modifiers, invalid procedure-diagnosis combinations) may represent enormous time cost for billing staff while the dollar recovery per appeal is modest.
- A lower volume of high-dollar clinical denials (medical necessity for inpatient admissions) represents a smaller queue but significantly larger recovery per resolved case.
- Run this analysis by payer, not just in aggregate. A denial pattern that looks manageable across your entire book may be heavily concentrated in two or three payers.
- That concentration matters for automation design — payer-specific rule libraries, portal login credentials, and appeal templates need to be built separately.

### The prioritization matrix:

- Plot your denial categories across two axes: automation feasibility (how rule-based and structured is the resolution path) and dollar impact (total at-risk revenue in this category over the past 12 months).
- The top-right quadrant high feasibility, high dollar impact is your Phase 1 automation target.
- Prior authorization denials and eligibility-related administrative denials almost always fall here.
- Clinical denials (medical necessity, level of care) typically have lower automation feasibility because resolution requires synthesizing clinical documentation against payer-specific medical policies.
- They belong in a later phase, once the RPA and ML layers are stable.

**Deliverable from this phase:** A ranked denial category list with dollar value, volume, and automation feasibility score for each category. This document governs your build sequence for the next 12 months.



### **Phase 3: RPA Layer - Configuration & Governance**

*RPA is the entry point for most teams. Configure it correctly the first time because bot debts are real and painful to unwind.*

#### The four workflows to automate first, in order:

- Eligibility verification at scheduling is the highest-leverage starting point because it prevents denials rather than recovering from them.
- Configure your bot to fire against the payer's 270/271 HIPAA eligibility transaction (or portal API where available) at the time of appointment creation, not the night before the visit.
- The earlier you surface a coverage problem, the more options the patient access team has to resolve it.
- Prior authorization status tracking is the second target.

- The bot should monitor open authorization requests across payer portals, pull status updates on a defined cadence, write the result back to the EHR encounter (via HL7 outbound interface or FHIR write endpoint), and flag any authorization that is expiring within 72 hours of a scheduled service date.
- Claim status monitoring handles the problem of claims sitting untracked in a submitted state.
- Configure the bot to query payer portals or the 277CA HIPAA transaction for all claims that have not received a determination within your defined SLA (commonly 14 days for commercial payers, 30 days for Medicaid).
- Aged claims trigger an automated follow-up action either a resubmission prompt or escalation to a human queue.
- ERA receipt and denial categorization is the fourth workflow. When an 835 ERA arrives, the bot reads the CARC and RARC codes, classifies the denial according to your Phase 2 taxonomy, calculates the payer-specific appeal deadline, and assigns the task to the appropriate work queue with the deadline and required action pre-populated.

#### **Governance requirements you cannot skip:**

- Every bot needs to run under a service account with credentials stored in a secret vault not hardcoded, not in a shared spreadsheet. Set a quarterly credential rotation schedule and build the vault integration before going live, not as a retrofit.
- Maintain a bot inventory document that is updated every time a new bot is deployed or modified.
- Shadow bots automation scripts built by end users outside the IT governance process are a HIPAA audit risk and a maintenance liability. The inventory requirement forces visibility.
- Define an exception escalation path for every bot before deployment. What happens when the payer portal is down?
- When a CAPTCHA breaks the session? When the ERA contains a reason code your taxonomy doesn't recognize?
- Every unhandled exception needs a defined human fallback, not a silent failure.

### ***Phase 4: ML Denial Prediction Model***

*This is where your system stops reacting to denials and starts preventing them.*

#### **Feature engineering what actually predicts denial:**

From your Phase 1 data, the features with the strongest predictive signal in denial models are: payer ID, CPT/ICD-10 code combination, DRG (for inpatient claims), place of service, attending physician NPI, days elapsed between prior authorization date and date of service, time since last claim for the same diagnosis under the same payer, and patient coverage type (commercial, Medicare Advantage, Medicaid).

Prior authorization status at time of claim submission is consistently one of the top predictors.

A claim submitted without a confirmed, non-expired authorization number, or with an authorization number that was obtained more than 30 days before the service date, has a materially higher denial probability across nearly all payer types.

### **Model selection and validation:**

Gradient boosting models (XGBoost or LightGBM) consistently outperform logistic regression on this feature set because of the categorical feature mix and the non-linear interactions between payer, procedure, and documentation status.

Do not use a single train/test split. Validate on out-of-time data — withhold your most recent six months from training and evaluate on them separately.

Payer policy shifts will cause a model trained only on older data to underperform on current claims.

Stratify your validation by payer. A single large commercial payer representing 40% of your volume can dominate your aggregate accuracy metrics while the model performs poorly on smaller payers that have different denial patterns.

Payer-stratified metrics surface that problem before it becomes a production issue.

### **Threshold calibration:**

Set your alert threshold based on your specific staffing and volume context, not a default. A threshold set too low floods reviewers with false positives and destroys adoption.

A threshold set too high misses the denials you built the model to catch. Run a cost analysis: what is the average administrative cost of a false positive (unnecessary human review) versus the average revenue loss of a false negative (missed denial)?

That ratio determines your calibration point.

### **Retraining cadence:**

Payer policies change frequently. Medicare Advantage plans in particular issue coverage determination updates that shift denial patterns within weeks of taking effect.

Build monthly drift detection into your pipeline monitor the distribution of incoming claim features and the model's confidence scores over time.

A meaningful shift in either signals that retraining is needed. Do not wait for revenue impact to tell you the model has degraded.

## ***Phase 5: NLP Appeal Engine***

*The highest-complexity layer. Scope it narrowly at first.*

**Do not start here.** Teams that lead with appeal automation before the RPA and ML layers are stable waste significant engineering effort automating the recovery of denials that a mature prevention layer would have stopped.

**Deploy Phases 3 and 4 first.** Appeal automation compounds value on top of a functioning prevention system.

### Scope for the first release:

Target one denial category with high volume, high appeal success rate, and a structured resolution path.

Administrative denials citing missing prior authorization documentation are ideal: the resolution action is consistent (attach the authorization to the appeal, submit through the payer's specified channel), the required documentation is retrievable from structured EHR fields, and appeal success rates for this category typically exceed 70% when the authorization was legitimately obtained.

### The generation pipeline:



The NLP layer needs to do four things in sequence. **First**, parse the denial reason from the ERA CARC/RARC codes and any accompanying payer correspondence.

**Second**, retrieve the relevant clinical documentation from the EHR using a FHIR DocumentReference query against the date of service.

**Third**, retrieve the payer's current published medical policy for the denied procedure this requires maintaining a payer policy document library, updated on a defined cadence (quarterly minimum, monthly for high-volume payers with frequent policy changes).

**Fourth**, generate a draft appeal in the payer's required format using the clinical documentation and policy as context.

### On using commercial LLMs:

If you are using an LLM API (OpenAI, Anthropic, or others) for the generation step, every request that includes patient data must flow through a BAA-covered deployment.

Confirm BAA coverage with your vendor before sending any PHI.

The generation step itself does not require fine-tuning passing clinical notes, the payer policy document, and a payer-specific appeal template as context at inference time produces adequate output quality for draft generation.

A human reviewer approves every draft before submission.

This is not optional in Phase 1 and should remain the standard even as the system matures.

### **Appeal deadline tracking:**

Build deadline logic into the categorization layer (Phase 3, RPA), not the appeal engine.

By the time a claim reaches the NLP appeal queue, its deadline should already be calculated, stored, and surfaced in the work queue interface.

The NLP layer's job is generation and routing, not deadline management.

## **Phase 6: Analytics, Feedback Loop & Iteration**

*A denial automation system that doesn't learn degrades. The analytics layer is what turns a deployment into a compounding asset.*

### **The four reports that matter operationally:**

Denial rate by payer, procedure code, and facility updated daily, not weekly. Revenue at risk by appeal deadline proximity segmented into less than 7 days, 8 to 30 days, and over 30 days.

Appeal success rate by denial reason code and payer. This is your system's report card and the leading indicator of where the NLP appeal engine needs improvement.

Root cause trending by upstream failure type which pre-service or pre-submission failures are generating the most denial dollars this period compared to last quarter.

### **The payer scorecard:**



Build a payer-specific view that tracks denial rate, average days to determination, appeal overturn rate, and response time to information requests per payer, over rolling 90-day and 12-month windows.

This report has two uses. Operationally, it identifies which payers require the most manual intervention and where both escalation thresholds need adjustment. Strategically, it is the quantitative foundation for contract negotiation conversations.

A payer whose denial rate is 3x the industry average for a specific DRG, with an 80% appeal overturn rate, is a payer you can have a data-backed conversation with at contract renewal.

### **Closing the ML feedback loop:**

Every claim that clears the denial-and-appeal cycle needs to feed back into the training dataset with its outcome labeled.

This is straightforward in principle and frequently broken in practice because the outcome data lives in a different system than the training pipeline.

Build the ETL process that joins ERA outcomes back to original claim submissions and reloads it into your feature store on a defined schedule.

Without this, your model trains on historical patterns that diverge further from current payer behavior with every passing month.

### **Iteration governance:**

Set a quarterly review cadence with a defined agenda: model performance metrics, both exception rates, appeal success rates by category, and the delta from baseline on total denial revenue.

The quarterly review is also when you decide which denial category moves from the human queue into the automated pipeline next.

This is the meeting that keeps the program progressing rather than plateauing after the initial deployment.

## **Common Failure Modes & How to Avoid Them**

**Skipping Phase 1.** Data quality problems discovered mid-implementation cost 3–5x what they would have cost to surface upfront. The audit is not optional.

**Automating too broadly in Phase 1.** Teams that attempt to automate 10 denial categories simultaneously in the first release typically deliver none of them well. Pick two. Ship them. Iterate.

**Bot credentials are stored insecurely.** Hardcoded payer portal credentials in RPA scripts are a compliance liability and a single point of failure. Vault integration is a prerequisite, not a later improvement.

**No human escalation path.** Every exception a bot cannot handle needs a defined next step before the bot goes live. Silent failures that drop claims from work queues are worse than no automation.

**Deploying the ML model without drift monitoring.** A model that performed at 82% accuracy at launch may be performing at 61% six months later as payer policies shift. Monthly drift detection is the difference between a model that stays useful and one that quietly erodes.

**Using the NLP appeal engine without BAA coverage.** PHI in appeal drafts is not a gray area. Confirm your LLM vendor's BAA status before any patient data touches the generation pipeline.

**No feedback loop into the ML model.** A denial prediction model that doesn't retrain on outcomes will diverge from reality. The ETL pipeline that closes this loop is as important as the model itself.

## Implementation Timeline Reference

Phase	Work	Typical Duration
1	Data readiness audit	2–3 weeks
2	Denial taxonomy & prioritization	1–2 weeks
3	RPA layer (2 workflows)	6–8 weeks
4	ML denial prediction (initial model)	8–12 weeks
5	NLP appeal engine (1 denial category)	10–14 weeks
6	Analytics dashboard & feedback loop	4–6 weeks (parallel to Phase 5)

Phases 3 and 4 can run in parallel after Phase 2 is complete. Phase 5 should not begin until Phase 3 is stable in production.

Total time to a functioning three-layer system: approximately 6–8 months for a team with dedicated engineering resources and clean historical data.

Add 6–12 weeks if the data readiness audit surfaces significant cleanup work.

## About Code-B

Code-B partners with healthcare software teams and revenue cycle leaders to architect, build, and integrate denial automation systems — from RPA configuration to custom ML model development and EHR integration.

Read the companion blog at [code-b.dev/blog/denial-automation-in-healthcare-apps](https://code-b.dev/blog/denial-automation-in-healthcare-apps) or get in touch at [manager@code-b.dev](mailto:manager@code-b.dev).

© 2025 Code-B Solutions Pvt Ltd. All rights reserved.