

Part 01

STACKAI WHITEPAPER

AI Agents for Healthcare Finance, Revenue Cycle, and Claims Operations

How providers and payers are deploying AI agents across denials, eligibility, documentation, and post-visit reporting to recover revenue that is currently written off, with the accuracy controls and audit trail that revenue leaders demand.

A practical guide for VPs of Revenue Cycle, CFOs, Directors of Coding and Compliance, and Billing Operations leaders.



PREPARED BY

Shani Fargun

VP Healthcare at StackAI

sfargun@stack-ai.com



HIPAA · SOC 2 TYPE II · ISO 27001 · GDPR

BAAs & DPAs with service providers

Executive summary

Revenue cycle is where healthcare loses money it has already earned. Claims are denied for fixable reasons and never reworked. Documentation misses a payer-specific element and a clean service becomes an uncollectible one sixty days later. Eligibility is wrong at the front door and the denial is baked in before the visit happens. Post-visit reporting to Medicaid is late or inaccurate and reimbursement slips.

The common thread is that this work is high-volume, rules-based, and judgment-heavy in ways that punish both full automation and pure manual effort. AI agents change the economics, but only if they are built the way revenue leaders require: built on your own denial history and payer rules, checked by an evaluator and a human, and granted autonomy progressively as measured accuracy earns it.

This whitepaper covers seven use cases proven across telehealth, home health, and multi-facility operators: post-claim denial analysis and resolution, closed-loop denial processing, pre-bill documentation compliance, eligibility and coverage capture, post-visit reporting and claim accuracy, finance close and reconciliation, and authorization lookup across state portals. In one home-health branch alone the denial workflow addresses more than 2,000 denied claims a week, work that previously consumed three to four case managers for roughly 16 hours a week. The design principle throughout: recommend first, approve with one click, and only automate the well-understood cases once accuracy clears the threshold you set.

2,000+

denied claims a week
now worked in a single
branch

**16
hrs/wk**

of manual denial
rework taken off 3–4
case managers

96%+

first-pass notes
compliance after pre-
bill review

55,000+

compliance audits run
a week in production

The revenue-leakage problem

Every revenue cycle leaks at four predictable points.

- Denials that never get worked. Denied claims pile up faster than staff can rework them, so a large share age out unrecovered. Each one has a specific, often fixable reason, but finding it, pulling the supporting documents, and filing the correction is slow manual work.
- Documentation that fails after the fact. A note is clinically sound but missing a payer-required element (a measurable goal, a time stamp, a medical-necessity justification). The claim goes out, gets denied two months later, and by then the encounter is impossible to reconstruct.
- Coverage errors at the front door. Wrong insurance on file, unverified eligibility, or a missing authorization means the denial is guaranteed before the visit even happens.
- Reporting that slips. For Medicaid-funded and facility-based operators, late or inaccurate post-visit reporting directly delays or reduces reimbursement.

This persists because the work does not scale with headcount, and the cost of an automated mistake in billing is high enough that most teams refuse to hand it to a black box. The fix is a governed agent with measured accuracy and a human in the loop.

Why revenue cycle needs different AI infrastructure

CHARACTERISTIC	WHY IT MATTERS FOR PLATFORM SELECTION
Accuracy is the whole game	A wrong code or a bad resubmission has financial and compliance consequences. The platform must support citations, evaluator scoring, and measurable accuracy before autonomy, not confident guessing.
Progressive autonomy	Revenue leaders will not flip a switch to full automation on day one. The platform has to support a staged path: recommend, approve, then auto-execute on well-understood cases, with the threshold under your control.
Built on your rules	Denials, payer requirements, and documentation standards are specific to your payers and your history. Agents must be built on your own knowledge base, and improve as approvals feed back into it.
Payer portal and clearinghouse reach	The work lives in Availity, payer portals, and clearinghouse feeds. The platform needs API and browser-automation reach into those systems.
Defensible audit trail	Every coding decision, correction, and resubmission has to be logged and explainable for payer audit and internal compliance.

The seven use cases

1

Post-Claim Denial Analysis and Resolution Recommendation

Pulls denied claims and their supporting documents, reads the denial reason against a knowledge base of prior resolutions, and recommends a specific corrective action.

2

Closed-Loop Denial Processing with Progressive Autonomy

Executes the approved resolution: files the correction, submits the appeal packet, or updates the record, with progressive autonomy.

3

Pre-Bill Documentation Compliance

Reviews every note against payer-specific requirements before submission and returns actionable feedback to the clinician.

4

Eligibility Verification and Coverage Capture

Verifies eligibility and extracts a structured coverage record from insurance cards at the front door.

5

Post-Visit Reporting and Claim Accuracy

Turns visit and call records into accurate, timely Medicaid and facility reporting.

6

Finance Close, Variance, and Reconciliation

Generates budget-versus-actuals executive summaries per facility and reconciles AP and vendor statements.

7

Authorization Lookup Across State Portals

Logs into state Medicaid and payer portals, captures authorization status and remaining units, and writes it back to the system of record.

1 Post-Claim Denial Analysis and Resolution Recommendation

THE PROBLEM

Denied claims arrive faster than staff can work them. Each denial requires pulling the claim, retrieving the specific supporting documents for that denial reason, understanding why it was denied, and deciding on a fix. At thousands of denials a week, most of the backlog ages out unrecovered.

HOW AN AI AGENT HANDLES IT

On a schedule, the agent pulls the claim-status feed from the payer portal (for example the 276 feed from Availity), and for every denied claim it retrieves the supporting documents required for that denial reason via API. It reads the denial reason against a curated knowledge base of prior denial reasons and their resolutions, and recommends a specific corrective action: a corrected code, missing documentation, an appeal letter, or a resubmission path. The recommendation is routed to a case manager for one-click approval through Teams, Outlook, or a form, and the approval decision feeds back into the knowledge base so the model improves.

WHY IT MATTERS

In one home-health branch this addresses more than 2,000 denied claims a week, work that three to four case managers spent roughly 16 hours a week on. Recovering even a modest share of previously-aged-out denials is direct, measurable revenue.

BUILT FOR	VP Revenue Cycle, Denials Management, CFO
KEY KPIs	Denials worked per week, recovery rate, revenue recovered, hours saved
TYPICAL IMPACT	Backlog worked systematically, recovery on previously-abandoned denials, case-manager time returned

2 Closed-Loop Denial Processing with Progressive Autonomy

THE PROBLEM

A recommendation is only half the job. Someone still has to file the corrected claim, generate and submit the appeal packet, or update the record, and route each action to the right system.

HOW AN AI AGENT HANDLES IT

This is the paired action layer. Once a case manager approves a recommendation (or once a case crosses the auto-approval threshold in later phases), the agent files the corrected claim, generates and submits the appeal packet, or updates the record, routing each action to the right downstream system. It is structured for progressive autonomy by design: it starts fully human-approved, then moves toward automatic execution on well-understood denial reason codes, with the accuracy threshold for autonomy set and controlled by you.

WHY IT MATTERS

Closing the loop is what turns "here is what to do" into recovered dollars without adding headcount, and the staged autonomy model is what makes revenue leaders comfortable letting it run.

BUILT FOR	Denials Management, Billing Operations, Compliance
KEY KPIS	Time-to-resolution, auto-resolution rate on eligible codes, appeal turnaround
TYPICAL IMPACT	Faster resolution, closed-loop recovery, autonomy that expands only as accuracy earns it

3 Pre-Bill Documentation Compliance

THE PROBLEM

A note that is clinically sound but missing a payer-required element becomes a denial after the claim is already out. By the time the denial arrives, the clinician has seen hundreds of patients and cannot reconstruct the encounter.

HOW AN AI AGENT HANDLES IT

An agent reviews every completed note against a rules engine built from your payer-specific documentation requirements, in real time or as a nightly batch. It checks for required structural elements (measurable goals, baseline and progress), billing alignment (does the documentation support the CPT code, is time documented for time-based codes), and compliance flags (consents, telehealth attestations, patient location). Notes that pass are cleared for billing. Notes that fail are returned to the clinician with specific, actionable feedback rather than a vague denial two months later.

WHY IT MATTERS

Catching a documentation gap before submission converts a future denial into a non-event. At scale, even a small first-pass-clean-claim improvement compounds into significant monthly revenue protected, and clinicians get immediate feedback that improves their documentation over time.

BUILT FOR	VP RCM, Coding and Documentation, Compliance, Clinical Operations
KEY KPIS	First-pass clean-claim rate, denial rate, days in A/R, clinician rework
TYPICAL IMPACT	Notes-compliance rate rises into the high 90s, denials drop proportionally, faster revenue recognition

4 Eligibility Verification and Coverage Capture

THE PROBLEM

Coverage errors at intake guarantee denials downstream. Manual eligibility checking and re-keying insurance cards is slow and error-prone, and "wrong insurance on file" is a leading root cause of avoidable denials.

HOW AN AI AGENT HANDLES IT

At the front door the agent verifies eligibility in real time and extracts a structured coverage record from an uploaded insurance card or PDF via OCR and extraction (payer, member ID, group, plan, effective date, copay, deductible), writing it to the chart so downstream eligibility, benefits, and billing flows work from clean data. A companion clinical-eligibility validator checks program-specific rules where they apply.

WHY IT MATTERS

It removes a manual data-entry step at intake and cuts the coverage errors that drive denials weeks later, protecting revenue at the point where it is cheapest to protect.

BUILT FOR	Patient Access, Revenue Integrity, Front-End Operations
KEY KPIS	Clean-coverage rate, front-end denial rate, intake data-entry time
TYPICAL IMPACT	Fewer coverage-driven denials, faster intake, clean data feeding the whole revenue cycle

5 Post-Visit Reporting and Claim Accuracy

THE PROBLEM

For Medicaid-funded and facility-based operators, reimbursement depends on accurate, timely post-visit reporting. Done manually across dozens of facilities, it is late, inconsistent, and a source of both revenue slippage and compliance risk.

HOW AN AI AGENT HANDLES IT

An agent turns visit and call records into structured, accurate post-visit and post-call reporting, applying the program's reporting rules, and routes exceptions for review. Deployed across a multi-facility footprint, it standardizes reporting so every facility reports the same way, on time, and flags the records that need a human before they go out.

WHY IT MATTERS

Accurate, timely reporting is the difference between full reimbursement and slippage, and across dozens of facilities the consistency alone materially reduces compliance exposure.

BUILT FOR	Revenue Cycle, Facility Operations, Compliance
KEY KPIS	Reporting timeliness, reporting accuracy, reimbursement realized
TYPICAL IMPACT	On-time, consistent reporting across facilities, protected reimbursement, reduced compliance risk

6 Finance Close, Variance, and Reconciliation

THE PROBLEM

Closing the books across a multi-facility network is a multi-day manual cycle: pulling GL data, computing variance, writing executive summaries by hand, and reconciling AP and vendor statements line by line.

HOW AN AI AGENT HANDLES IT

Once AP and AR close each period, an agent connects to the ERP, pulls the GL data, computes variance against budget and prior period, and generates an executive-summary narrative tailored to the reader (facility administrator, regional VP, CFO), routing site-specific summaries to the right recipient. Companion flows produce standalone variance reports for controllers and flag unusual GL activity. On the AP side, a reconciliation agent ingests supplier statements each month, matches them line by line against internal records, and flags missing invoices, duplicate payments, and aging discrepancies, with row-level permissioning so each analyst sees only the vendors and facilities they are entitled to.

WHY IT MATTERS

It turns a multi-day close into same-day, consistent output across facilities, and gives leadership comparable financials without waiting on manual assembly.

BUILT FOR	CFO, Controllers, Facility Administrators, AP
KEY KPIS	Days to close, reconciliation exceptions caught, reporting consistency
TYPICAL IMPACT	Multi-day close becomes same-day, consistent per-facility reporting, faster exception review

7 Authorization Lookup Across State Portals

THE PROBLEM

Care coordinators spend hours a day logging into state Medicaid and payer portals, one patient at a time, to check authorization status, remaining service units, and expiration dates. It is slow, it is easy to miss a lapse, and a missed authorization means unpaid care already delivered.

HOW AN AI AGENT HANDLES IT

A browser-automation agent logs into the portals on the operator's behalf, navigates to the authorization screens, captures current status, units remaining, expiration dates, and pending actions, and returns a structured record for each patient. It runs the loop across the full patient list on a schedule and writes the extracted data back into the system of record (for example AlayaCare), so the care-management team works from one source of truth instead of portal after portal. Because this is all PHI, the workflow runs inside the operator's single-tenant environment with encrypted, bot-only access to portal credentials.

WHY IT MATTERS

It removes hours per day of manual portal work per coordinator and catches authorization lapses before care is delivered against an expired auth, which is one of the most preventable causes of unpaid work in home care.


BUILT FOR	VP Operations, Care Management, Revenue Cycle
KEY KPIS	Coordinator hours saved, authorization-lapse rate, unpaid-care avoidance
TYPICAL IMPACT	Hours per coordinator per day recovered, fewer lapsed authorizations, single source of truth

SEE IT IN PRODUCTION

Three of these workflows, built and running on StackAI today.

Pre-submission denial prevention

Eligibility, modifier, diagnosis-to-procedure, and documentation checks run in parallel, feed a risk assessment, and generate a specific fix list plus a Teams alert to the billing team before the claim is ever filed.



app.stack-ai.com



Claim Denial Prevention

What this does
Submit claim details and the related note. We'll run pre-submission checks (eligibility, coding, documentation, payer rules) and return a **risk score** and **fix list** before you file.

Uploads (at least one)

- Clinical note (progress note / therapy note / operative note)
- Optional: orders, labs/imaging, referral, prior auth confirmation

Claim Data

Patient Eligibility Data

Clinical Documentation

Outputs

Denial Prevention Report

Result History

Submit

The surface billers use: submit claim data and the note, get a risk score and fix list back before filing.

Pre-bill documentation compliance

A clinical note is checked against payer-specific documentation requirements and returns a compliance score, a pass-or-fail determination, and specific fixes with example language the clinician can use.

StackAI · workflow canvas

Clinical Note

The basic input node for the flow.

Session Date: 2024-08-12

Payment Code: Age 34

Duration: 45 minutes

Modality: Telehealth (video)

Patient reports continued anxiety and difficulty sleeping, and discussed recent stressors at work and practical challenges managing. Patient engaged well. Plan: continue weekly sessions, adjust to individualized therapy.

Diagnosis: F41.1 Generalized Anxiety Disorder

Clinical Note Document

Enables user to upload file (e.g. PDF, Word, Excel etc.)

Expected on input

Next step: file uploaded

CPT Code

The basic input node for the flow.

20234

Payer

The basic input node for the flow.

Aetna

Session Duration (minutes)

The basic input node for the flow.

45

Modality

The basic input node for the flow.

telehealth

Patient State

The basic input node for the flow.

OK

Diagnosis

The basic input node for the flow.

F41.1 Generalized Anxiety Disorder

Compliance Reviewer

Anthropic Agent with tool calling

Choose LLM Model

Knowledge Sources

Tools

MCQ

ADD MCP

Compliance Report

Generate a rich template from inputs

Clinical Documentation Compliance Review

CPT: [CPT Code] | Payer: [Payer] | Duration: [Session Duration (minutes)] | Modality: [Modality] | Date: [Session Date]

View Results

Compliance Report C

Output the result of nodes

View Results

The build: the note plus claim metadata feed a compliance reviewer that produces a structured report.

app.stack-ai.com

Check your note before you bill.

Instructions

Paste or upload a clinical note and I'll review it against payer-specific documentation requirements before claim submission. You'll get a compliance score, pass/fail determination, and specific fixes for any gaps with example language you can use.

Fill in:

Your note (paste text or upload)

CPT code, payer, session duration

Modality, patient state, diagnosis

Takes about 10 seconds. Catches the issues that cause denials 60 days later.

Clinical Note

Type something...

Clinical Note Document

Choose or drag and drop here

Maximum file size of 20Mb

CPT Code

Type something...

Payer

Type something...

Session Duration (minutes)

Type something...

Modality

Type something...

Patient State

Type something...

Diagnosis

Type something...

The surface: check a note before you bill, in about ten seconds, and catch what would deny sixty days later.

Financial statement reconciliation

Two financial statements are compared into a side-by-side balance-sheet table, with a summary of the material differences and a drafted email of the findings ready to send.

StackAI · template

Financial Statements Reconciliation Agent

This agent compares two financial statements and presents a markdown table with a side-by-side comparison of the balance sheet.

AI Agents in this Workflow

OpenAI 1

You are an AI assistant specialized in financial reporting. - Be concise and professional. - Highlight only material differences (ignore...

Anthropic

You are an AI assistant designed to compare two financial statements. Your response should include a Markdown table that presents a side-...

Integrations

AI

Gmail

OpenAI

Back to templates

Use template

```
graph LR; U1[Upload Financial Statement (Excel)] --> AI[AI]; U2[Upload Financial Statement (PDF)] --> AI; AI --> E[Extract & Compare Financial Sections]; E --> V[Visual Comparison Table]; V --> R[Register Financial Impact]; R --> S[Summary of Differences & Insights]; R --> D[Draft an email of Differences & Insights];
```

The build: upload two statements, extract and compare the sections, then summarize and route the differences.

Common questions, answered

Revenue leaders push hardest on accuracy and control. Here are the questions we hear, and the answers.

Q1 "How do I know it will not code something wrong or file a bad claim?"

Three controls, and none of them is "trust the model." First, agents run on your own denial history and payer rules and cite their sources, and can be instructed not to act unless confident. Second, a second evaluator model scores the output against your rubric, and multi-agent decomposition keeps each step narrow, which is where accuracy comes from. Third, a human approves every consequential action until you decide otherwise. You cannot always eliminate model error entirely, but unlike a human, the agent logs why it reached each decision, so every action is explainable and reviewable.

Q2 "So is this fully autonomous, or not?"

It runs autonomously only where measured accuracy has earned it. The denial workflow is built for progressive autonomy: it starts as recommend-and-approve, moves to one-click approval, and only auto-executes on well-understood denial reason codes once measured accuracy clears the threshold you set. You control the threshold and which code categories qualify. Nothing flips to autonomous on day one.

Q3 "Our denials live in Availity and payer portals. Can it reach them?"

Yes. The denial workflow pulls the claim-status feed (for example the 276 from Availity) and retrieves supporting documents via API, and where a payer system has no usable API we use browser automation. The recommendation and approval loop runs through Teams, Outlook, or a form, wherever your case managers already work.

Q4 "Can it stand up to a payer audit?"

That is the design intent. Every coding decision, correction, and resubmission is logged with its inputs, the source rule or prior resolution it relied on, the output, the approver, and a timestamp, and it is exportable. The human-in-the-loop approvals create a defensible decision trail, and the documentation-compliance workflow itself improves the defensibility of the underlying notes.

Q5 "How do we prove ROI before committing?"

A predictable monthly subscription, not usage-based, with a scoped proof of value on your highest-leakage workflow. We measure recovery and hours saved on your own volumes: a single branch working 2,000-plus weekly denials, or a documentation-compliance lift on your actual payer mix, produces a concrete payback number before any annual commitment.

How these use cases compound

- 01 **Eligibility and coverage capture** stop denials at the cheapest point to stop them, the front door.
- 02 **Documentation compliance** stops the next tier of denials before claims ever go out.
- 03 **Denial analysis and closed-loop processing** recover the denials that still happen, and feed every resolution back into the knowledge base so the system gets smarter.
- 04 **Post-visit reporting** protects the reimbursement that depends on accurate, timely filing.

Each layer reduces the load on the next, and the approvals feeding back into the knowledge base are what let autonomy safely expand over time.

The value case

METRIC	BEFORE → AFTER AGENTS		IMPACT
Denials worked	before	Backlog ages out after Systematically worked	Recovered revenue
First-pass clean-claim rate	before	Baseline after High-90s notes compliance	Denials drop proportionally
Front-end denial rate	before	Coverage errors common after Reduced	Protected reimbursement
Denial rework labor	before	3 to 4 FTEs, ~16 hrs/week per branch after Reallocated to exceptions	Skilled hours returned
Reporting timeliness	before	Late, inconsistent after On-time, consistent	Reimbursement protected

Why StackAI for revenue cycle

StackAI is the AI transformation platform for healthcare: a HIPAA- and GDPR-compliant operating system for enterprise agentic AI, built by MIT PhDs and deployed across more than 200 organizations in regulated industries.



Accuracy controls

Citations, evaluator scoring, and measured accuracy before autonomy, not confident guessing.



Progressive autonomy

Recommend, approve, then auto-execute on eligible codes, with the threshold under your control.



Built on your data, self-improving

Agents run on your denial history and payer rules, and improve as approvals feed back in.



Payer-system reach

API and browser automation into Availity, payer portals, and clearinghouse feeds.



Audit trail

Every decision logged, explainable, and exportable for payer audit and compliance.



Predictable pricing

Flat monthly subscription with ROI proven on your own denial and documentation volumes.

The bottom line

Revenue cycle loses money that has already been earned, at four predictable leak points, because the work does not scale with headcount and the cost of an automated mistake is too high to hand to a black box. Governed AI agents, built on your rules, checked by an evaluator and a human, and granted autonomy only as accuracy earns it, recover that revenue with a trail a payer auditor will accept.

Want to see the recovery math on your denials?

Book a working session with the StackAI healthcare team to scope a proof of value on your highest-leakage workflow, with recovery and hours saved measured on your own numbers.