Part 04 STACKAI WHITEPAPER

AI Agents for Home-Based, Senior, and Community Care

How home-health operators, senior-care networks, and residential and I/DD providers are deploying AI agents to survive the paperwork: authorizations, case-management packets, survey readiness, and month-end close.

A practical guide for COOs, CFOs, VPs of Operations and Compliance, and Revenue Cycle leaders in home health, hospice, senior living, and community-based care.

PREPARED BY

Shani Fargun

VP Healthcare at StackAI

sfargun@stack-ai.com

Executive summary

Home-based, senior, and community care runs on the thinnest margins in healthcare and the heaviest administrative load. The work that keeps the lights on is unglamorous and relentless: logging into state Medicaid and payer portals to check authorizations, processing case-management packets from a dozen different companies with a dozen different rulebooks, reconciling shift notes against incident reports before a state survey, and closing the books across facilities every month.

None of this is clinical, and almost none of it requires the EHR. It is exactly the kind of high-volume, rules-based, cross-system work that governed AI agents do well. This whitepaper lays out five use cases proven in production across a home-health operator, a multi-facility senior-care network, and a residential and I/DD provider: authorization lookup across state portals, case-management packet processing, residential quality and incident-reconciliation audits, finance close and reconciliation, and a cited regulatory policy assistant.

In production these are workhorses. One residential provider runs roughly 28,000 staff-level analyses and 13,000 individual-level compliance reports per 180 days. One home-health operator processes 200 to 300 case-management packets a month at about an hour each today, and faces more than 2,000 denied claims a week in a single branch. The agents that absorb this load do not replace care staff; they give a thin-margin operator back the hours and the audit trail it cannot otherwise afford.

02 THE PROBLEM

The paperwork-survival problem

Providers in this segment operate under a specific kind of pressure. Reimbursement is largely Medicaid and Medicare, so rates are fixed and margins are slim. Regulation is intense and enforced by periodic state survey, where a documentation gap can mean a citation, a corrective action plan, or a licensing risk. And the operational surface area is enormous: many patients, many payers, many state portals, many case-management partners, many facilities, each with its own forms and rules.

The result is that a large share of the workforce spends its day on manual, repetitive, error-prone administrative work: logging into portals one patient at a time, reading packets against rulebooks, hunting for discrepancies between systems, and rebuilding the same reports every month. It is expensive, it does not scale, and it is precisely where a small documentation error becomes a denied claim or a survey finding weeks later.

This is the highest-yield place in healthcare to deploy AI agents, because the work is high-volume, rules-based, cross-system, and almost never requires touching a clinical system of record.

03 PLATFORM FIT

Why this segment needs different AI infrastructure

CHARACTERISTIC	WHY IT MATTERS FOR PLATFORM SELECTION
PHI plus portal automation	Much of the work means logging into state Medicaid and payer portals with credentials, against patient data. That demands single-tenant deployment, encrypted bot-only credential handling, and browser automation that runs inside your environment.
Many partners, many rulebooks	Packets arrive from a dozen-plus case-management companies, each with its own forms and required fields. The platform has to classify the source and apply the right rulebook, not assume one standard.
Survey and audit readiness	Every automated action needs a defensible, exportable trail, because a state surveyor will ask you to prove it.
Thin margins demand hard ROI	These operators cannot fund science projects. The platform has to prove time and revenue recovery on your own volumes during a short, predictable-cost pilot.
Small or non-standard systems	Home-care systems like AlayaCare, eVero, Monday.com, and SharePoint, not Epic. The platform has to orchestrate across exactly this kind of mixed, smaller-vendor stack.

04 THE FIVE USE CASES

The five use cases

Finance close and reconciliation leads the set: it turns a multi-day, multi-facility close into same-day, comparable output for facility administrators, regional VPs, and the CFO. Four more workflows run in production across a home-health operator, a senior-care network, and a residential and I/DD provider, taking on the portals, packets, and audits that surround the books.

1

Finance Close, Variance, and Reconciliation

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

2

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.

3

Case-Management Packet Processing

Classifies each incoming authorization packet by sending company, applies that company's rulebook, and flags discrepancies before it is actioned.

4

Residential Quality and Incident-Reconciliation Audit

Audits behavioral plans and tasks against regulatory criteria, and reconciles shift notes against incident reports, at site, individual, and staff level.

5

Regulatory Policy Assistant

Cited, always-current answers on HCBS, OPWDD, ICF/IID, and state licensing rules.

1 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

2 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

3 Case-Management Packet Processing

THE PROBLEM

Authorization packets arrive from 12 to 20 different case-management companies, each with its own document standards, forms, and required fields. Staff process them by hand against the sending company's specific rules, at roughly an hour per packet, 200 to 300 packets a month.

HOW AN AI AGENT HANDLES IT

An agent monitors the case-management inbox, ingests mixed-format inputs (PDFs, ZIPs, Word documents, images), classifies the sending company, retrieves that company's processing rulebook and prior documentation from SharePoint, looks up the client in the system of record, and runs a discrepancy check comparing the packet against the patient record. It flags anomalies (mismatched authorization dates, service-hour discrepancies, missing physician orders, wrong plan of care) and generates a recommendation of what has to change before the packet can be actioned.

WHY IT MATTERS

It compresses an hour of skilled manual review into a fast exception check, and it standardizes accuracy across partners so a packet from any of a dozen companies is checked the same way, every time.

BUILT FOR	Intake and authorization teams, Compliance, Operations
KEY KPIs	Minutes per packet, discrepancy-catch rate, rework and resubmission rate
TYPICAL IMPACT	An hour per packet becomes a fast exception review, consistent multi-partner accuracy

4 Residential Quality and Incident-Reconciliation Audit

THE PROBLEM

Residential and I/DD providers are judged on survey readiness. They must prove that each individual's behavioral plan and recorded tasks meet regulatory criteria, and that shift notes and behavior documentation reconcile with the incident reports on file, across every site. Done by hand, this is impossible to do comprehensively, so gaps surface at the worst possible time: during the survey.

HOW AN AI AGENT HANDLES IT

An audit agent compares each individual's behavioral plan and recorded tasks from the EHR against regulatory criteria, then reconciles shift notes and behavior documentation against incident reports. It runs on a schedule across every site and produces a site-level quality report, an individual-level report, and a staff-level analysis identifying which staff are driving compliance gaps. A sibling incident-reconciliation flow closes the long-standing gap between the shift-notes system and the incident-reporting system. In production this runs at roughly 28,000 staff-level analyses and 13,000 individual-level reports per 180 days: the workhorse the operator relies on to survive state licensing audits.

WHY IT MATTERS

It converts survey readiness from a frantic pre-survey scramble into a continuous, documented state, and it pinpoints exactly where and with whom the gaps are, so remediation is targeted.

BUILT FOR	Compliance, Quality, Residential Operations
KEY KPIs	Documentation-gap rate, reconciliation coverage, survey-finding rate
TYPICAL IMPACT	Continuous, comprehensive audit coverage, targeted remediation, defensible survey trail

5 Regulatory Policy Assistant

THE PROBLEM

Clinical, care, and compliance staff need answers on state and federal rules governing residential and community programs (HCBS, OPWDD, ICF/IID, state licensing standards), and those rules change.

HOW AN AI AGENT HANDLES IT

A chat-based assistant answers natural-language policy questions with verbatim citations back to the underlying regulation and internal policy. A regulation-parser subflow keeps the knowledge base in sync with the latest published state and federal updates, and companion flows crawl agency sites for changes, so answers never drift from current rule.

WHY IT MATTERS

It gives frontline staff a fast, cited, always-current answer instead of a guess, which is both a productivity tool and a compliance control in a survey-driven environment.

BUILT FOR	Compliance, Quality, Care Management, Program Directors
KEY KPIs	Time-to-answer, policy currency, answer consistency
TYPICAL IMPACT	Fast, cited, current answers, reduced compliance risk from stale guidance

SEE IT IN PRODUCTION

The finance-close workflow, built and running on StackAI today.

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, the same shape that powers the multi-facility month-end close.

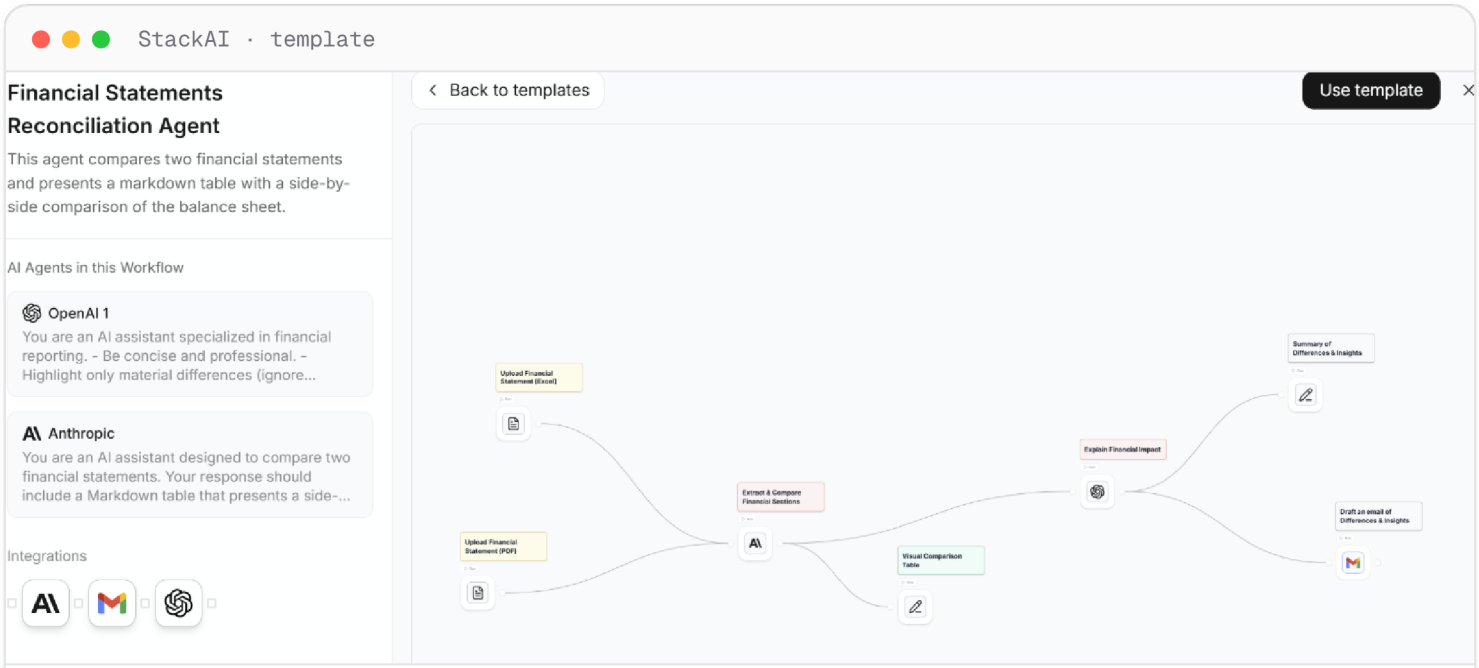

The build: upload two statements, extract and compare the sections, then summarize and route the differences.

05 — COMMON QUESTIONS, ANSWERED

Common questions, answered

The objections in this segment center on trust, stack fit, survey defensibility, and hard payback. Here are straight answers.

Q1

“We would be giving an agent our state portal logins, against patient data. Is that safe?”

Yes, and it is designed for exactly this. Authorization-lookup and packet workflows run inside your single-tenant environment, where PHI never leaves your control. Portal credentials are encrypted and handled bot-only, the browser automation runs in your environment, and every action is logged. StackAI is HIPAA compliant and SOC 2 Type II with BAAs in place, and for this segment we default to single-tenant precisely because of the portal-and-PHI combination.

Q2

“Our systems are AlayaCare, eVero, Monday.com, and SharePoint, not Epic. Does that matter?”

It works in your favor. StackAI is an orchestration layer built to connect across exactly this kind of mixed, smaller-vendor stack, and these workflows do not require a clinical EHR integration at all. Where a system has an API we use it, and where it does not (older portals, for instance) we use browser automation.

Q3

“Will this hold up in a state survey?”

That is the point of the audit trail. Every automated action is logged with its inputs, the source rule or document, the output, the reviewer, and a timestamp, and it is exportable to your tracker of choice. The residential quality audit is specifically built to produce site, individual, and staff-level evidence, and human-in-the-loop review covers anything consequential, so you can show a surveyor exactly what happened and why.

Q4 “Margins are thin. How do we know this pays off before we commit?”

A predictable monthly subscription, not usage-based, so costs never surprise you. Engagements start with a three-month proof of value on your highest-pain workflows, and we measure the ROI on your own volumes (hours saved on portals and packets, denials recovered, close time reduced) before any annual commitment. Given that a single branch can see 2,000-plus denied claims a week and packets take an hour each, the payback math is usually fast and concrete.

Q5 “Can it really process packets from all our different case-management partners?”

Yes. The agent classifies the sending company first, then applies that company's specific rulebook, so it handles a dozen-plus partners with different forms and required fields rather than assuming one standard. New partners are added by loading their rulebook, not by rebuilding the workflow.

06 COMPOUNDING VALUE

How these use cases compound

01 **Authorization lookup and packet processing** attack the two biggest sources of unpaid work and manual hours at the front of the revenue process.

02 **The quality and incident audit** keeps you survey-ready continuously instead of scrambling, and pinpoints where to remediate.

03 **Finance close and reconciliation** give leadership fast, comparable numbers across facilities.

04 **The policy assistant** grounds all of it in current, cited regulation, so the same source of truth feeds authorizations, audits, and staff answers.

Together they give a thin-margin operator back hours, revenue, and an audit trail it otherwise cannot afford.

07 — THE $ VALUE

The value case

METRIC	BEFORE → AFTER AGENTS	IMPACT
Portal authorization work	before — Hours per coordinator per day after — Scheduled, automated	Coordinator time recovered
Case-management packet	before — ~1 hour each, 200 to 300 / month after — Fast exception review	Skilled hours returned
Survey documentation coverage	before — Sampled, pre-survey scramble after — Continuous, comprehensive	Defensible readiness
Month-end close	before — Multi-day manual cycle after — Same-day, consistent	Faster, comparable reporting
Denied-claim exposure	before — Thousands per week per branch after — Systematically worked	Recovered reimbursement

08 WHY STACKAI

Why StackAI for home-based and community care

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.

Single-tenant for PHI

Portal automation and packet processing run in your isolated environment, with encrypted bot-only credentials.

Browser automation

Reaches state Medicaid and payer portals that have no usable API.

Mixed-stack connectivity

Orchestrates AlayaCare, eVero, Monday.com, SharePoint, and your ERP, no Epic required.

Survey-ready audit trail

Every action logged and exportable, with human-in-the-loop on consequential steps.

No-code builder

Compliance and operations staff add new partner rulebooks and reports without engineering.

Predictable pricing

Flat monthly subscription with ROI measured on your own volumes during the pilot.

09 THE BOTTOM LINE

The bottom line

The work that keeps a home-care, senior-care, or residential operator solvent is the paperwork, and it is drowning the workforce. Governed AI agents can take on the portals, the packets, the audits, and the close, inside your own environment, with an audit trail a surveyor will accept, at a cost that pays for itself in recovered hours and revenue.

Want to see this on your volumes?

Book a working session with the StackAI healthcare team to scope a four-to-six week single-tenant pilot on your highest-pain workflow, with ROI measured on your own numbers.

StackAI