

Part 03

STACKAI WHITEPAPER

AI Agents for Health Systems, Academic Medical Centers, and Community-Based Clinics

How hospitals, faculty practice plans, and community health centers are deploying governed AI agents around the EHR, not through it, to relieve administrative load without touching the system of record.

A practical guide for CIOs, CMIOs, CTOs, VPs of Operations, and Compliance and Revenue leaders evaluating enterprise AI agent platforms.



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

Health systems, academic medical centers, and community-based clinics share a hard truth: the administrative and operational burden around care has grown faster than the workforce, and the EHR is not the place to solve it. Epic, Oracle Health, and the rest are systems of record, not automation engines, and every integration into them is a slow, IT-gated, instance-specific project.

The opportunity is to deploy governed AI agents around that system of record: on the access center, the patient portal, the knowledge base, the intake queue, and the eligibility and prior-auth front door. These are high-volume, high-friction workflows that rarely require deep read or write access to Epic, and where an agent with the right guardrails delivers measurable relief in weeks rather than quarters.

This whitepaper lays out five use cases proven in enterprise and academic settings, including a clinical triage and evaluation pipeline built at a leading academic medical center. It addresses head-on the concerns every health-system buyer raises: EHR integration reality, data residency and on-premise control, clinical governance and hallucination, and why a governed orchestration layer succeeds where Microsoft Copilot Studio runs out of room.

The systems that adopt this posture will free clinical and operational staff for the work that requires them, with an auditable trail for every automated action.

The administrative-burden problem

The published research is consistent: clinicians and staff spend a large and growing share of their time on documentation, messaging, prior authorization, and coordination rather than on patients. In a health system that burden is distributed across the access center, the coding and revenue-integrity teams, care coordination, compliance, and the clinicians themselves.

The instinct is to ask the EHR vendor to fix it. But three realities get in the way:

- The EHR is a system of record, optimized for storing and retrieving the chart, not for orchestrating cross-system workflows or making judgment calls.
- Every integration is an instance-specific IT project. Even two organizations on the same Epic build have different endpoints, preferences, and data structures, so there is no single reusable connector, and lead times of twelve weeks or more are normal.
- Vendor AI add-ons are tied to a single ecosystem and a single model, with limited configurability and no independent governance layer.

The systems getting meaningful leverage deploy a governed agent layer around the EHR instead of waiting for it to become an automation platform.

Why health systems need different AI infrastructure

CHARACTERISTIC	WHY IT MATTERS FOR PLATFORM SELECTION
The EHR is not an automation engine	The highest-ROI agent workflows live around the chart, not inside it: access center, portal, intake, knowledge base, eligibility. A platform that assumes deep EHR write access will sit in your IT queue.
Data residency is non-negotiable	AMCs and many systems require that PHI never leave their environment. The platform must run single-tenant or fully on-premise, with your own models and a specific cloud region, so residency commitments are enforceable rather than aspirational.
Clinical governance is the gating factor	Nothing consequential ships without review. Human-in-the-loop, citations, evaluator scoring, and a complete audit trail are prerequisites, not features.
No vendor conflict of interest	An orchestration layer is not trying to sell you more cloud, more licenses, or more of one model family. That neutrality matters when the workflow spans Epic, Microsoft, Google, and homegrown systems.
Buy the platform, own the workflows	Subject-matter experts in revenue integrity, access, and compliance should be able to build and change agents themselves, without a professional-services engagement for every edit.

The five use cases

Five workflows recur across enterprise and academic settings, chosen because they deliver measurable relief around the EHR without waiting on a deep integration.

1

Clinical Triage and Intake Routing (with Evaluator)

Classifies inbound patient messages and intake responses by urgency and clinical category, routes to the right pathway, and scores every triage against a clinical rubric.

2

Access Center and Call QA and Compliance

Near-total audit coverage of access-center calls and messages against HIPAA, consent, script, and quality standards.

3

Staff Knowledge and Policy Assistant

Sourced, cited answers for clinical and operational staff across SOPs, protocols, formulary, and state and federal regulation.

4

Patient Portal Message Management

Classifies portal messages and drafts responses for staff review, cutting the in-basket load.

5

Eligibility, Prior-Auth, and Coverage Front Door

Verifies eligibility, extracts coverage from insurance cards, and assembles prior-auth packets before the visit.

1 Clinical Triage and Intake Routing (with Evaluator)

THE PROBLEM

Inbound patient messages and intake responses arrive faster than staff can triage them, and the cost of getting urgency wrong is high. Manual triage is slow, inconsistent between staff, and invisible to quality review.

HOW AN AI AGENT HANDLES IT

An agent ingests an inbound message or intake response, classifies urgency and clinical category, and routes it to the appropriate care pathway. An evaluator loop scores every triage decision against the organization's own clinical rubric, so accuracy is measured continuously rather than assumed. This is the exact shape of pipeline built at a leading academic medical center as its first evaluation of agentic AI in clinical intake, deliberately deployed in a demo-and-evaluator posture first, with a documented path to expand into intake automation and provider-facing triage aids as confidence grows.

WHY IT MATTERS

Triage is where AI in a clinical setting has to earn trust, so it is the right place to prove the governance model: measured accuracy, cited reasoning, and a human reviewing anything the agent is not fully confident about, before autonomy is ever extended.

BUILT FOR	CMIO, VP Clinical Operations, Access Center leadership, Quality
KEY KPIS	Triage accuracy against rubric, response time, escalation appropriateness
TYPICAL IMPACT	Faster routing, measured and auditable triage quality, a governed path from evaluation to production

2 Access Center and Call QA and Compliance

THE PROBLEM

Access centers and patient-facing teams handle enormous call and message volume subject to HIPAA, consent requirements, and internal quality standards. Manual QA reaches only a few percent of interactions, leaving the vast majority unreviewed.

HOW AN AI AGENT HANDLES IT

An agent reviews every call transcript, chat, and message against a structured rubric: identity verification and minimum-necessary standards, consent and disclosure, script and protocol adherence, resolution and tone, and escalation appropriateness, plus sentiment and safety-signal detection. Results roll up by agent, team, and channel, and flag failing interactions for human QA. The same shape extends to voice-agent calls, which are summarized into structured intent and QA records.

WHY IT MATTERS

Coverage moves from a sampled few percent to near-total, at a cost that scales sub-linearly with volume, turning QA from a spot-check into a continuous, audit-ready trail and reducing regulatory exposure.

BUILT FOR	VP Patient Access, Compliance, Quality, Revenue Integrity
KEY KPIs	Percent of interactions audited, issue-detection rate, escalation accuracy
TYPICAL IMPACT	Near-total audit coverage, earlier detection of systemic issues, defensible audit trail

3 Staff Knowledge and Policy Assistant

THE PROBLEM

Frontline clinical and operational staff burn time hunting through SOPs, protocols, formulary, payer rules, and constantly changing state and federal regulation. Answers are inconsistent, and tribal knowledge walks out the door with turnover.

HOW AN AI AGENT HANDLES IT

A retrieval-based assistant answers natural-language questions across the internal knowledge base with verbatim citations back to the source document and section. A regulation-parser subflow keeps the knowledge base in sync with the latest published state and federal updates, so answers do not drift from current rule. Companion collector workflows capture new answers from senior staff and add them to the base, so coverage improves without a dedicated knowledge-management team.

WHY IT MATTERS

For a community-based clinic or an academic practice navigating HRSA, HCBS, state Medicaid, and payer-specific rules, a cited, always-current policy assistant is both a productivity tool and a compliance control.

BUILT FOR	Clinical operations, compliance, revenue integrity, care coordination
KEY KPIS	Time-to-answer, answer consistency, policy-currency, onboarding ramp
TYPICAL IMPACT	Faster, consistent, cited answers, reduced dependence on tribal knowledge

4 Patient Portal Message Management

THE PROBLEM

The patient portal in-basket has become one of the heaviest hidden burdens in a health system, and much of it is administrative rather than clinical.

HOW AN AI AGENT HANDLES IT

A classifier reads every inbound portal message and routes it (scheduling, refill, billing, clinical, escalation). A draft-reply agent pulls the relevant context and drafts a response for the appropriate staff member or provider to review and send, with clinical guardrails and escalation criteria built in. Administrative messages resolve fast, clinical ones arrive at the provider pre-summarized.

WHY IT MATTERS

Portal burden is a documented driver of clinician burnout. Cutting the routine load, while keeping a human on anything clinical, returns time to the people who need it most.

BUILT FOR	Ambulatory operations, clinic managers, CMIO
KEY KPIS	In-basket volume per provider, response time, staff time per message
TYPICAL IMPACT	Routine messages deflected or pre-drafted, clinical messages arrive summarized, faster turnaround

5 Eligibility, Prior-Auth, and Coverage Front Door

THE PROBLEM

Coverage errors and missing authorizations discovered after the visit are a leading cause of denials and write-offs, and for community-based clinics and safety-net providers the margin to absorb them is thin.

HOW AN AI AGENT HANDLES IT

At the front door, an agent verifies eligibility in real time and extracts a structured coverage record from an uploaded insurance card or PDF (payer, member ID, group, plan, copay, deductible), writing it to the chart so downstream billing works from clean data. For services that need authorization, the agent assembles the prior-auth packet from the record against payer requirements and routes it for review. A companion clinical-eligibility validator checks program-specific rules and returns a pass, flag, or hold with the rule cited.

WHY IT MATTERS

Catching coverage and authorization issues before the visit converts a downstream denial into a non-event, and for a grant- or Medicaid-funded organization, that revenue integrity is existential.

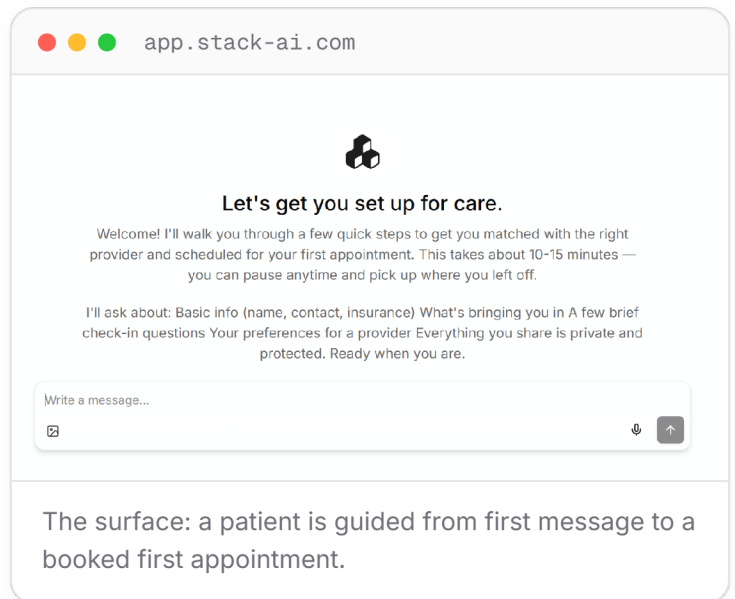

BUILT FOR	Revenue Integrity, Patient Access, community-based clinic finance leadership
KEY KPIS	Clean-coverage rate, prior-auth turnaround, front-end denial rate
TYPICAL IMPACT	Fewer denials from bad coverage data, faster authorizations, protected reimbursement

SEE IT IN PRODUCTION

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

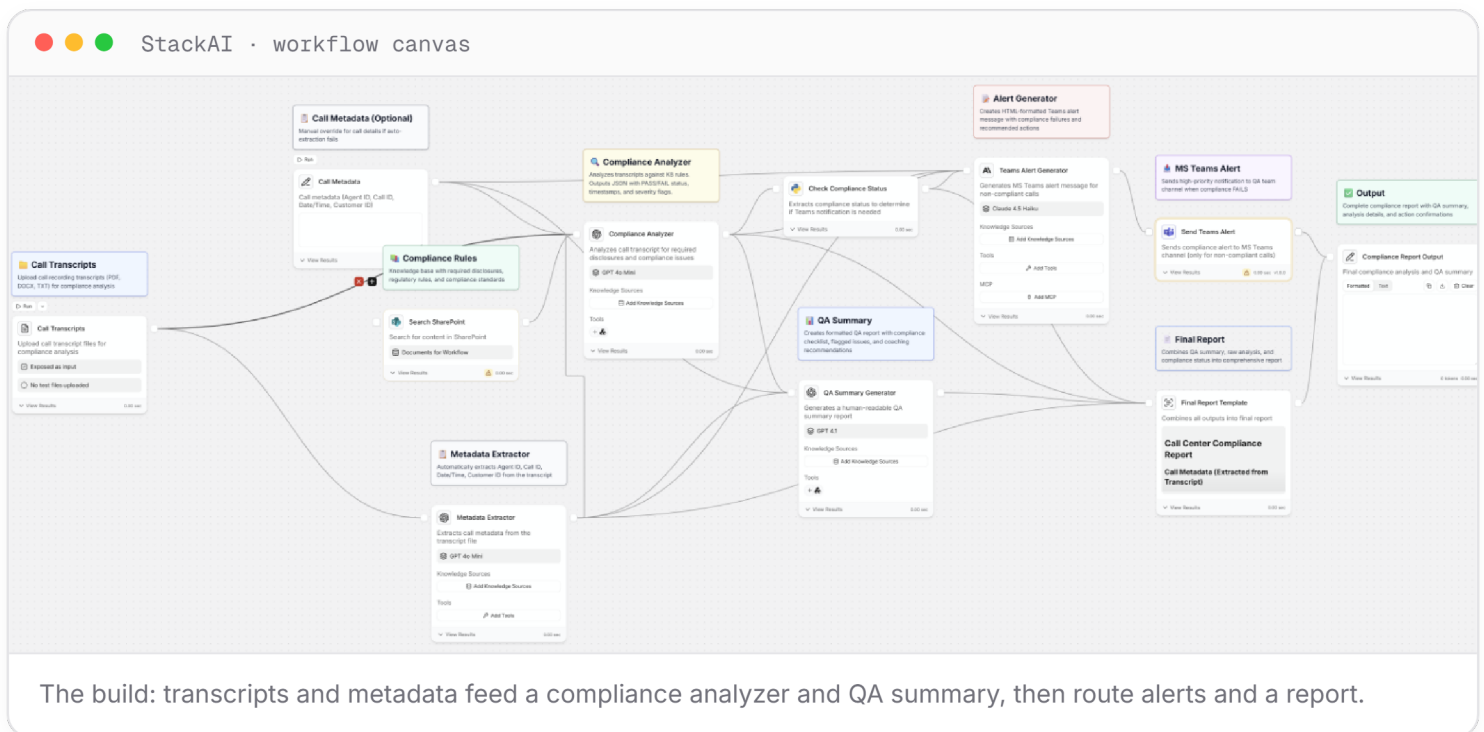
Intelligent intake and triage

An intake orchestrator runs eligibility, provider matching, and scheduling as subflow tools, turning a patient's first message into a structured intake and a booked appointment.



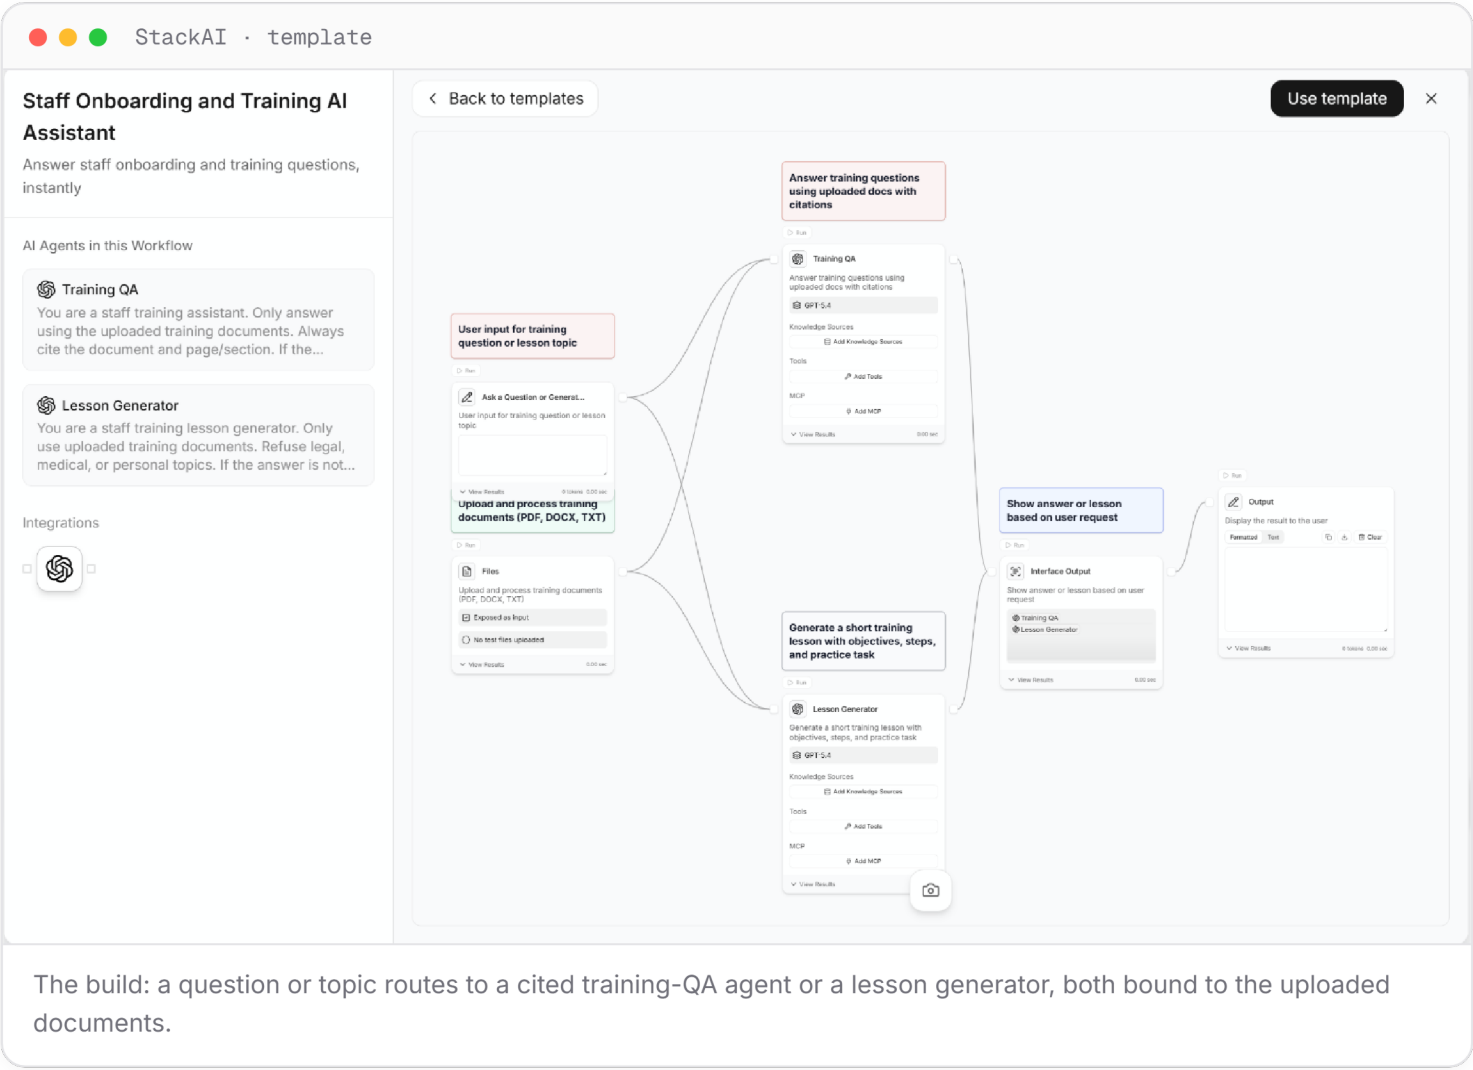
Access-center call QA and compliance

Call transcripts are analyzed against a compliance rulebook, producing a QA summary, a pass-or-fail status, a Teams alert on non-compliant calls, and a full compliance report.



Staff knowledge and training assistant

A training assistant answers staff questions only from the uploaded documents, always citing the source, and a lesson generator turns a topic into a short lesson with objectives, steps, and a practice task.



Common questions, answered

Health-system buyers ask hard questions. Here are the ones we hear most, with straight answers.

Q1 “We are on Epic. Can you even integrate, and how long will it take?”

Direct Epic integration is possible but slow and instance-specific: every Epic build is configured differently, and IT lead times of twelve weeks or more are normal. So we design around it. Most of the five use cases above do not need deep read or write access to Epic at all, because they run on the access center, the portal, the knowledge base, and the intake queue. Where chart data is genuinely needed, we prefer the lightest path that works: targeted FHIR calls for specific data elements, an existing MCP server or App Orchard integration, an intermediary such as Redox, or RPA, rather than a full custom Epic build. StackAI is the orchestration layer on top of your stack, working alongside the EHR.

Q2 “Our PHI cannot leave our environment. Can you run on-premise?”

Yes, and for AMCs this is usually the deciding factor. StackAI deploys single-tenant or fully on-premise, in your own cloud, in the region you specify, with your own local models. PHI never leaves your control, which is precisely why this posture is chosen by the most residency-sensitive academic and enterprise customers. A specific-region deployment commitment is enforceable, not aspirational, because the software runs where you put it.

Q3 “How do we govern this clinically, and prevent hallucination?”

Governance is built into the product. Agents answer only from approved, cited sources and can be told not to respond unless confident. An evaluator model scores outputs against your clinical rubric. Every clinical or revenue workflow ships with human review, approval, and escalation, and autonomy is only ever extended once measured accuracy clears a threshold you set. Every action is logged with inputs, model, output, reviewer, and timestamp, which is the audit trail your compliance and quality teams need.

Q4 "We are a Microsoft shop. Why not Copilot Studio?"

Copilot works inside the Microsoft ecosystem, but health-system teams consistently hit the same walls: connecting a working knowledge base is difficult, reaching non-Microsoft or clinical data requires extra connectors and glue, custom interfaces and complex logic are limited, and standing it up meaningfully needs a dedicated team. A neutral orchestration layer has no incentive to sell you more cloud or seats, spans Epic, Microsoft, Google, and homegrown systems equally, and ships governance and model fallback out of the box.

Q5 "How is this priced, and how do we prove ROI before committing?"







A predictable monthly subscription, not usage-based. Engagements start with a three-month proof of value: up to three production workflows, builder seats, and a weekly build-and-support session with an AI strategist and forward-deployed engineer. We measure ROI on your own volumes during the POV, so the business case is proven before any annual or on-premise commitment.

How these use cases compound

- 01 **Access-center and QA monitoring** is the governance foundation that lets you deploy clinical-adjacent agents with confidence.
- 02 **The knowledge and policy assistant** makes every other workflow more accurate, because the same cited, current source of truth grounds triage, portal replies, and eligibility rules.
- 03 **Triage and portal management** relieve the two heaviest sources of frontline and clinician burden.
- 04 **The eligibility and prior-auth front door** protects revenue and feeds clean data into everything downstream.

The result is administrative relief around the EHR, delivered in weeks, with a measured and auditable trail for every automated action.

The value case

METRIC	BEFORE → AFTER AGENTS		IMPACT
QA audit coverage	before 	2 to 5 percent	Continuous, audit-ready trail
	after 	Near-total	
Portal in-basket per provider	before 	Rising	Time returned to clinicians
	after 	Reduced routine load	
Triage quality visibility	before 	Sampled or none	Measured, defensible clinical quality
	after 	Scored on every case	
Front-end denial rate	before 	Baseline	Protected reimbursement
	after 	Reduced	
Policy answer time	before 	Minutes of searching	Faster, consistent decisions
	after 	Seconds, cited	

Why StackAI for health systems

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.



Deployment control

Secure deployment on your cloud or StackAI's, in the region you choose, with your own models. PHI never leaves your environment.



Around the EHR, not through it

Orchestrates access center, portal, knowledge base, and eligibility without requiring deep Epic write access.



Clinical governance

Human-in-the-loop, citations, evaluator scoring, and complete audit logging on every workflow.



Vendor neutrality

Model-agnostic with fallback, no conflict of interest with your cloud or EHR vendor.



No-code builder

Revenue-integrity, access, and compliance experts build and change agents without an engineering ticket.



Universal connectivity

500+ integrations across the Microsoft, Google, Epic, and homegrown systems you already run.

The bottom line

The administrative burden around care has outgrown the workforce, and the EHR was never designed to solve it. Governed AI agents deployed around the system of record can, in weeks rather than quarters, and without asking your IT team to fight the integration queue first.

**Want to start deploying secure
AI agents at scale, with
consistency?**

Book a working session with the StackAI healthcare team to scope a three-month proof of value on the workflows that do not require deep EHR access, with a clear on-premise or single-tenant path for the ones that touch PHI.



© 2026 StackAI. Confidential. Prepared for health system, academic medical center, and community-based clinic leaders.