

Your Health System Is Already Running on Agents.

Is Anyone Actually Managing Them?

By **Abhinav Shashank**
April 14, 2026

Since the beginning of this year, I've come across more than 100 companies launching or expanding AI agents designed to run inside health system workflows.

Scheduling agents.

Prior auth agents.

Clinical documentation agents.

Revenue cycle copilots.

Patient communication agents.

Hospital operations command centers.

Some are from the electronic health records (EHR) vendors themselves, some from cloud platforms, some from startups nobody had heard of 18 months ago. All sold separately. All designed to plug into the same operational stack.



Nobody is asking the obvious question: When agents from five, six, seven different vendors are touching the same patient workflow, who is responsible for what happens between them?

THE TOOLS WORK, WHAT'S THE PROBLEM?

A 2026 survey of 120 health systems found that 75% are now running at least one AI application. Half are running three or more. Add in the AI features embedded natively in the EHR and the math gets bigger fast. The major EHR platforms are each shipping over a hundred AI features this year, from ambient charting to autonomous coding to revenue cycle copilots.

The tools work. That's not the problem.

The problem is that nobody is managing the space between them.


The first wave is over.

For the last two years, health systems have been in the adoption phase. An AI governance committee would evaluate a tool, run a pilot, maybe do a security review, and approve it. The American Medical Association (AMA) published an eight-step intake toolkit. The Joint Commission and CHAI released guidance in late 2025. NIST has a risk management framework.

All of it was built for a world where you adopt AI tools one at a time. That world is gone.

What you're left with is fragmented AI. A prior auth tool here, an ambient documentation layer there, a population health platform from a third vendor. Each evaluated in isolation. None designed to operate together. BCG found that 74% of organizations haven't seen tangible value from their AI investments yet, despite years of experimentation. That's not a capability problem. It's a coordination problem.

We're now in the second wave. The question isn't whether to adopt AI. It's how to run an organization where dozens of autonomous agents are making decisions across scheduling, documentation, coding, prior authorization, population health and care management, often without talking to each other, and without any single person watching the whole picture.



Becker's Hospital Review called this the "accountability vacuum." Individual agents perform correctly in isolation. The failures show up in the interactions, and they're hard to spot because the workflow keeps running. A cardiac marker flagged by one agent doesn't get passed to a recommendation agent reading imaging data. The system doesn't crash. It just makes a confident, wrong decision.

And it gets worse. Research from late 2025 documented how multiple AI agents, each designed to offer clinical second opinions, can create "consensus pressure" that overrides correct guideline-based reasoning. The agents agree with each other because they're drawing from similar data. The resulting false consensus looks, to both the system and the clinician, like validation. Nobody questions it because nobody is watching the interactions between agents. They're watching the agents individually.

THREE THINGS THAT NEED TO EXIST — AND MOSTLY DON'T

If you're running a health system with three or more AI tools in production, here's what I think you need. Not next year. Now.

First, an owner. Not a committee. A person. Someone with operational authority over the agent layer, the same way you have a CMIO for clinical informatics or a VP of Revenue Cycle. That role needs to run something: a standing function with use case prioritization, governance across vendors, and accountability for outcomes. Some systems are getting this right. UC San Diego Health has a Chief Health AI Officer. Mayo Clinic and Cleveland Clinic have appointed Chief AI Officers. Sutter Health, UCSF, Hackensack Meridian, Children's National, all of them created

dedicated AI leadership roles in the last year. But most health systems are still governing AI through their IT committee structure. That structure was built for procurement decisions. Not production management.

Second, a shared data layer. This one goes deeper than most people realize. It's not just that agents conflict without shared data. Even isolated point solutions struggle to produce good results without one. An ambient documentation tool that only sees what happens in the visit, without longitudinal history, produces notes that miss critical context. A coding agent that reads claims data but can't see the clinical record

generates confident, wrong codes. Today, 70-85% of AI pilot failures trace back to data quality and infrastructure gaps, not to the models themselves. The AI isn't broken. The data underneath it is fragmented.

Now multiply that across three, four, five vendors. Your EHR vendor's agents see EHR data. Your RCM vendor's agents see claims data. Your population health vendor's agents see registry data. They're all touching the same patient, often on the same day, each working with a partial picture.

This is the real interoperability problem of 2026. Not FHIR compliance. Not API standards. Whether the agents running inside your organization, individually and collectively, share enough context to do their jobs.

Third, orchestration as infrastructure. Not monitoring after the fact but architecting for coordination from the start. Agents need shared context, defined handoff protocols, and a common view of the patient record. When a prior auth agent triggers a denial, a downstream agent should know it automatically. When

a scheduling agent and a care management agent both touch the same encounter, they should be working from the same facts. Network security isn't a quarterly review. It's a continuous operational function. Agent orchestration needs to work the same way.

The EHR vendors would argue they solve this. They don't. Not fully.

Some EHRs now offer agents for scribing, revenue cycle, patient scheduling, and even agent-builder frameworks that share a data layer and governance structure. Within their own ecosystem, the orchestration works.

But everything stops at the vendor's edge. The prior auth automation tool from a different vendor, the population health platform flagging care gaps, the patient engagement solution sending outreach, those are separate data pipelines and separate accountability structures. The typical health system isn't running one vendor's stack. It's running four or five. The question of who governs the whole is unanswered.

WHAT MANAGING AGENTS ACTUALLY LOOKS LIKE

The health systems getting this right are treating agent governance the way they treat supply chain management. Not every vendor decision is made in isolation. Someone is responsible for the whole.

Practically, that means putting agent orchestration under operations, not just IT. Treating agent interactions as a production environment that needs active management, not a policy question that needs a quarterly meeting. And asking a different question during every vendor evaluation: not "what does this agent do?" but "how does it interact with the agents we already have?"

The technology is moving fast. The organizational design is standing still. We have tools from 2026 running inside org charts from 2015.

It's the problem that drove us to build Gravity at Innovaccer, the coordination layer underneath all of our workflows and agents. A shared data foundation that every agent draws from, and an orchestration layer that manages handoffs between them regardless of which vendor built them, and an interoperability layer that connects them back to your systems of record. The agents can only work together if something is designed to make them work together.

The question isn't whether your health system is running AI agents. It already is. The question is who's accountable when they don't work together.

AUTHOR



Abhinav Shashank is the cofounder and CEO of Innovaccer, a San Francisco-based healthcare AI company on a mission to transform how care is delivered and experienced. He is championing a bold vision for Autonomous Healthcare — a future where intelligence is embedded into every workflow, friction is engineered out of care delivery, and technology works quietly in the background so clinicians and patients can focus on what matters most.