THE REVENUE LAYER

When Every Agent Succeeds, Why Is the Patient Still Waiting?

December 2026 · By Jen Viens

HLTH 2026 made it clear that we are no longer talking about AI in healthcare as a possibility, a pilot or something we should consider. The market has moved from pilots to production agents, and the real question now is not whether you can buy another agent but whether you can govern and orchestrate them. The patient is still waiting because the control plane that coordinates agents across organizations, the foundation for collaboration across agents with different properties, is still being built.

Every platform now coordinates the agents inside its own walls, and the patient's journey crosses those walls. Most of the energy at HLTH was on what agents can now do, not on how health systems and plans will govern fleets of agents from multiple vendors. That is clearly next, and it creates a clear opening for anyone selling into health systems and plans.

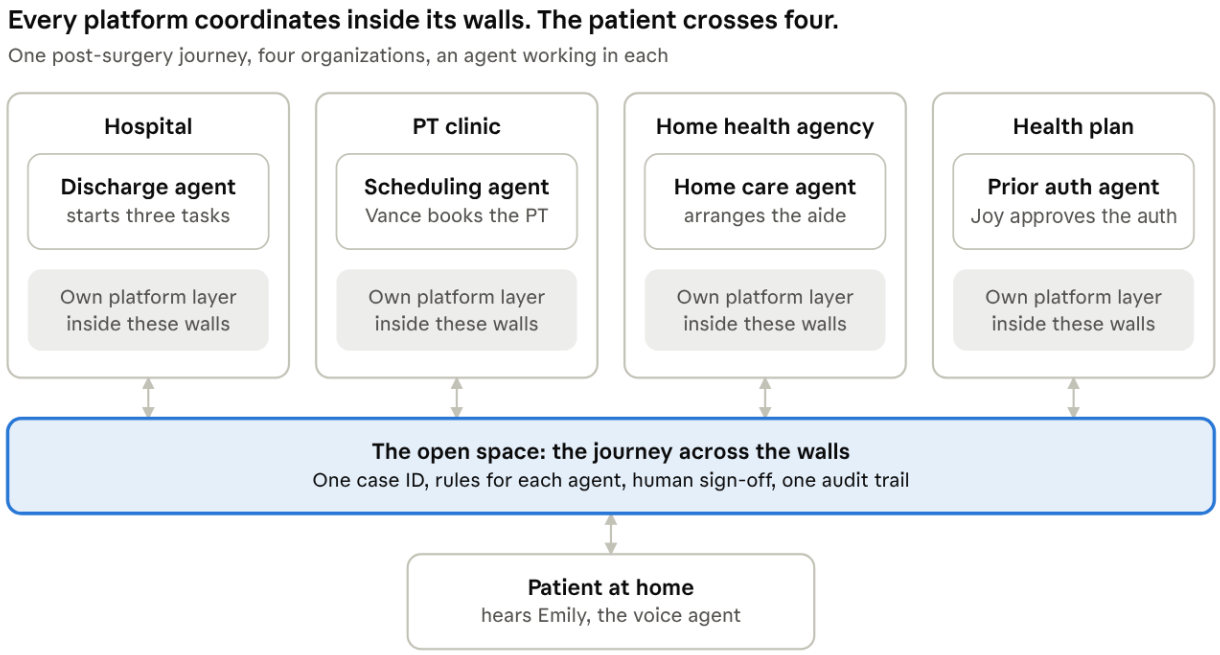
Agents on both sides

Agents already work both sides of the transaction. On the provider side, voice agents have placed more than 10 million automated calls between providers and insurers, and the CEO of the company placing them calls the clash "bot wars" (SiliconANGLE). On the payer side, the largest US insurer has AI agents calling doctors' offices to schedule appointments (Bloomberg). Health plans have committed to answering, in real time in 2027, at least 80% of electronic prior authorization approvals submitted with complete documentation (BCBSA). For anyone selling into this market, that changes the buyer's question. It is no longer only what your agent does, but how it works with every other agent touching the same patient.

One patient, four organizations

A patient has surgery. The hospital's discharge agent must work with outpatient care for PT, a home health aide and the health plan's prior authorization: three more bots and three tasks, coordinated securely on one plane. Picture it from the patient's side. Emily, the voice agent, calls: she is coordinating with Vance on the outpatient PT, and the prior authorization was approved by Joy. All of the agents are coordinating care in real time, on a securely managed control plane.

What does success look like? The agent drops off, or a human is contacted or takes over; the appointment gets booked, and the patient is engaged and feels comforted. Voice-enabled agents are already doing parts of this today. Hackensack Meridian's voice agent calls patients within 72 hours of discharge, flags those at higher risk of readmission and alerts the clinical team to follow up (Chief Healthcare Executive). Tampa General's voice agent cut call abandonment from 34% to 14.9% and increased scheduled appointments by 21% (Becker's).



With one layer across the walls, the patient hears it from Emily instead of carrying the handoffs.

The team underneath

This is enterprise team management, and it needs the same foundation as a coordinated care team: training, security, audit and management. Leadership and management set a clear foundation and one playbook, so everyone collaborates toward the same goal and a known, clear, understandable outcome. For this to succeed, the collaboration and the outreach have to produce the result, including how it is delivered by voice and how it is managed on the agentic control plane. It goes back to setting the goals and driving the execution.

In practice, that layer is not another agent. It is the referee and the record keeper: one case ID that follows the patient, written rules for what each agent may see and do, a person's sign-off on clinical content, and one audit trail. Every agent has a named owner, a level of autonomy and a way to pause it.

When the agents disagree

The work is just sent back to a clinician as more hours of work, but that is not an answer. So how do we manage it: as a queue, with a human, or with a super bot?

The floor is already set. Under the health plans' 2026 commitments, AI may approve, and a licensed and qualified clinician reviews any medical necessity denial (AHIP and BCBSA). Seven states passed laws this year limiting AI's role in coverage decisions (Becker's Payer). Physicians already complete about 40 prior authorizations a week (AMA).

If it is mandated, that is the answer: the decision stays with a clinician. However, the doctors and nurses are overtaxed. Outside prior authorization, we have to have a data log of how to handle these

disagreements, so doctors and nurses can do what they do best: treat.

The AMA's new governance framework for agentic AI puts four questions to any agent that takes on clinical functions. What function is it performing? What shows it can perform it, and what happens if it fails? What is it permitted to do, and under what conditions? Who is accountable (AMA)? Expect those questions from the buyer's clinical and compliance teams, and have the answers ready before they ask.

How I would position it

Picture a large regional health system in a very progressive spot. It is running three different agents: one for voice outreach, one for prior authorization and one for post-discharge follow-up. Each works in isolation. None of them talk to each other cleanly, write-backs to the EHR are inconsistent, and no one can answer which agent did what, under whose authority and with what patient data. There is no flow, so the patient is the only coordinator.

Instead of selling another agent, I would run a short agentic journey design. Map the real discharge-to-home-health journey, define the target agent sequence, set clear human gates, and design a thin orchestration layer with a single case ID and policy controls. Think of it as a basic data orchestration map. The first pilot would focus only on progress-note completeness and appointment coordination after orthopedic discharge, measured on follow-up completion and on an audit trail that compliance can actually use. If it works, it becomes the blueprint for taking the same governed pattern to prior authorization and member outreach.

The lesson is that you win by helping the health system design the map and the rules for an agentic workflow.

Price the journey, not the agent

Agents are often bought one at a time as point solutions, each with its own pilot and contract, and too many of those pilots never connect or never make it to close. For multi-agent systems in healthcare, the consumption model works best when it is aligned to the journey.

- Sell the governed journey or outcome chain, not isolated agents.
- Price and package around the orchestration and policy layer and the highest-value journeys: post-acute activation, prior authorization and outreach, discharge to follow-up.
- Let best-of-breed agents, voice and prior authorization among them, live under the same plane, so consumption grows there.
- Make expansion natural. Once the control plane and the audit pattern exist, the next journey or agent is incremental, not a new science project.

In a multi-agent flow, consumption follows the flow, not the agent count: how many agents touch the case, how much context they pass to each other, and how far they reach for the data. Anthropic's engineers found that multi-agent systems use about 15 times the tokens of a chat and pay off only where the task is valuable enough (Anthropic). That is why inference closer to the data layer is where the value is: less data moved, less context passed between agents, and faster answers. It also means less patient data in each agent's hands, which is the minimum-necessary rule. Model the flow, then price it. Minutes

and tokens still meter the work, but shift the conversation from how many agents to how the flow works and what it does for the patient: governed journeys, not wholesale agents.

Predictability matters to buyers: in a Google Cloud report, 40% of healthcare organizations cited unpredictable operating costs as a reason AI initiatives could be paused or abandoned (via Healthcare IT Today). Outcome pricing works where the outcome can be counted and credited, such as a completed booking or a completed post-op contact. Otherwise, price the journey and keep the outcome as the measure.

What sales leaders should do now

First, build the framework with the client before you add another agent. Define the outcomes, responsibilities and rules of engagement. Just as a care team needs training, leadership, security protocols and clear accountability, so does a team of AI agents. Who owns the outcome? What information can each agent access? Who makes decisions, and when must a person step in? For the seller, these questions are the discovery call.

Second, start with one complete patient journey. Take a patient moving from surgery to discharge, home health, outpatient physical therapy and follow-up care. Map every handoff across the hospital, physician, health plan and patient. Give the participating agents a governed way to exchange authorized information, coordinate next steps and escalate exceptions. Make one journey work before expanding to others. This is also the land. In every account, know whose platform you will sit beside, because every EHR and cloud now claims the orchestration layer.

Third, measure the outcome, not the activity. How many calls an agent completes or how quickly it responds matters far less than whether the patient's care is actually coordinated. Was the appointment scheduled? Was the authorization resolved? Did someone intervene when needed? Most importantly, did the patient understand what was happening and feel supported throughout? Voice is key here, because it is how the patient hears that their care is coordinated. These are the measures the buyer renews on, so sell to them from the first meeting.

The bigger picture

We have spent years building AI that can perform tasks, and now we need to build the management and coordination that allows those capabilities to work together.

I've led large organizations and complex partner ecosystems. One lesson holds across all of them: individual performance does not guarantee a successful outcome. Clear leadership, shared objectives, defined responsibilities and accountability are what turn individual capabilities into an effective team, and AI agents are no different. Authority gets earned the same way: an agent takes on more only after it proves itself, which is how one governed journey becomes the next.

Production isn't simply the day an agent goes live. It's the day the patient no longer has to coordinate the system themselves.

Jen Viens is Founder of The Revenue Layer, providing fractional CRO leadership and commercial advisory to enterprises, early-stage technology companies and investors.