

Agentic AI in healthcare is too big for most to build alone

Should we build it ourselves? For payers contemplating agentic AI in healthcare solutions, that may be the wrong question.

The [healthcare AI market](#) was estimated to be more than \$36 billion in 2025, and [92% of healthcare executives](#) believe AI will provide a competitive edge. But enthusiasm has outpaced honest accounting of what it costs to build and operate agentic AI in production.

The gap between what works in a demo and what works at scale remains the single biggest barrier to value and crossing that gap requires far more than most payers anticipate.

Agentic AI in healthcare is an operating layer, not just a feature

Most payers approaching agentic AI think of it as a tool that slots into existing workflows. That framing underestimates the problem.

More than just generating output, agentic AI systems act across workflows: coordinating appointments, pulling real-time eligibility data, and managing multistep member journeys from first inquiry to confirmed care. For example, a chatbot that recommends a doctor isn't the same thing as a system that books the appointment, verifies insurance coverage, and follows up to collect valuable Patient-Reported Outcome Measures (PROMs) and Patient-Reported Experience Measures (PREMs).

That distinction changes the scope of what “building” means. A tool requires development. An operating system requires architecture, governance, ongoing maintenance, real-time data infrastructure and possibly data modernization, clinical validation, and human oversight running simultaneously.

Frontier AI companies like Anthropic are rapidly improving the underlying intelligence behind AI. But in healthcare, the model itself is only part of the solution. Pager Health SM sits on top of that foundation, turning AI capabilities into real-world healthcare experiences through the workflows, governance, integrations, and member support systems needed to make AI actually work in payer environments.

The hard truth: most internal builds fail

The gap between a convincing internal demo and a production-grade agentic AI is where most initiatives quietly collapse. Three failure patterns explain why.

The demo trap

Internal teams can build impressive AI demos quickly with clean data, scripted scenarios, and controlled inputs. Then the system meets production: messy data, edge cases, users who behave unpredictably, and errors with real clinical and financial consequences. A successful demo proves that AI can work in ideal conditions. It doesn't prove it should be deployed, or that it will perform at scale.

A successful demo is often mistaken for proof of readiness, when it's just the beginning of a much harder build.

The illusion of control

Building feels strategic. But [Gartner](#) projects that by 2027, 60% of organizations will fail to realize the anticipated value of their AI use cases. The failure mode isn't usually technical, but organizational: timelines stretch, costs mount, and systems require constant rework as data changes and models drift.

What looks like control is actually exposure to risk, cost, and delay that compounds over time.

The talent and complexity gap

Production-grade agentic AI for care navigation requires a stack most payers are not structured to operate: AI/ML engineering, clinical validation, real-time systems integration across provider directories and scheduling platforms, ongoing model evaluation, and compliance oversight. Each is a discipline in its own right. Assembling and retaining that talent is a different challenge than most build-vs-buy analysis account for.

The [costs of implementing AI](#) don't stop at launch. Annual maintenance, covering model monitoring, retraining, security, and compliance, typically runs 15%–30% of the original build cost.

The problems payers can’t afford to “figure out later”

Even payers with the resources to build internally face structural challenges that can't be deferred. These are foundational requirements that determine whether the system is fit to deploy at all.

1. Safety isn’t optional

Safety can't be bolted on after deployment. Agentic [AI in healthcare needs guardrails](#), audit trails, human-in-the-loop escalation, bias monitoring, hallucination controls, and clear accountability from day one.

The market is moving in this direction. The National Academy of Medicine released an [AI Code of Conduct framework](#) in 2025 to support AI that performs accurately, safely, reliably, and ethically in health and medicine. Without adequate safeguards, agentic AI fails, creates liability, and erodes member trust.

2. Real-time data is the foundation

Agentic systems that act on stale or inaccurate data underperform and mislead. A member directed to a provider who isn't accepting patients, or toward a service their plan doesn't cover, experiences a failure of care access.

Without live provider validation, eligibility checks, and scheduling integration, the autonomy agentic AI promise becomes an illusion.

3. Workflows are the product

Care navigation is a journey: inquiry, eligibility, provider matching, scheduling, and follow up. A system that answers questions fluently but can't complete the task hasn't solved the problem. The measure of success is a confirmed appointment with the highest quality, in-network provider that results in the best possible care outcome for the member.

Designing for task completion requires orchestration across digital and staff-supported channels; a level of complexity most internally built pilots never reach.

Deploy the system. Don’t rebuild it from scratch.

Leading plans are winning by getting proven systems into production faster, with less risk, and with accountability built in. Partnering with a purpose-built platform delivers:

- **Speed to value:** months to deployment, not years, with infrastructure already tested in complex healthcare environments
- **Built-in safety:** clinical guardrails, human escalation, security, governance, and continuous monitoring embedded from day one
- **Real-time data infrastructure:** validated integrations across provider directories, eligibility, and scheduling, all without EHR-level lift
- **Shared accountability:** Incentives aligned to outcomes, not just implementation milestones

The goal isn't to hand off AI strategy. A partner who brings proven infrastructure, shared accountability, and the ability to move fast without sacrificing governance allows payers to deploy production-grade capability now, while internal teams stay focused on what payers are built to do.

The winners will be pragmatic, not ambitious

The future of agentic AI in healthcare belongs to those who know what not to build.

The payers that come out ahead won't be those that built the most. They'll be those that deployed responsibly, integrated deeply, and scaled with governance structures to sustain it.

Payers that reach production faster are already demonstrating the outcomes that improve Star performance, deflect call center volume at scale, and strengthen member engagement.

Explore how [Pager Health Provider Navigator](#) delivers agentic AI in healthcare with clinical guardrails, real-time data connections, security, and human escalation built into every workflow, without the build timeline.

Connect with Pager Health to explore how AI-powered orchestration can help your plan simplify member navigation, improve engagement, and deliver measurable ROI.

First Name *

Last Name *

Business Email *

Organization *

Role

How can we help? *

Message

Submit

Solutions	Customers	Resources	About	Contact
Pager HealthSM Navigator	For Payors	Blog	Company	Contact Us
ReallyWellSM –	For Providers	White Papers	Careers	
Member Login	For Employers	Case Studies	Press	
Enterprise 360				

Pager HealthSM is committed to making our website accessible to everyone, including individuals with disabilities. We strive to conform to the Web Content Accessibility Guidelines (WCAG) 2.1 Level AA or higher and applicable provisions of the Americans with Disabilities Act (ADA) and Section 508 of the Rehabilitation Act.

We regularly review and update our website to maintain and improve the accessibility of our digital content and services.

This statement is reviewed and updated periodically to reflect our ongoing commitment to accessibility. See our full [Accessibility Statement](#) for more information.

