How agentic AI is reshaping health plan operations

AUTHORS

Sally Kim
Director, Health Plan Research, Advisory Board

Scott Brown
SVP, Advisory Services, Optum Advisory

Troy Anderson, Ph.D.
VP, Advisory Services, Optum Advisory

PUBLISHED

August 2026

Understanding the landscape

What is agentic AI?

Agentic AI refers to systems that do more than analyze data or generate content. These systems can plan, make decisions, and take actions toward a defined goal, often using multiple tools and workflows with limited human oversight.

In contrast, **traditional AI** is typically rules-based or predictive. It performs well for narrowly defined tasks such as classifying claims, flagging potential fraud, or forecasting costs. **Generative AI** expands on that ability to create new outputs, including drafting member communications, summarizing clinical notes, or answering questions in natural language. However, generative AI still depends on human prompts and direction at each step.

Agentic AI goes further by connecting these capabilities into end-to-end workflows. Agentic AI can identify a problem, decide which data to review, generate analyses or communications, trigger operational actions, and adjust based on results. While each agentic AI tool is optimized for a narrow function, their combined use unlocks broader, more adaptive AI capabilities. For example, an agentic system could detect rising prior authorization delays, investigate contributing factors, initiate outreach, and track whether the intervention improved turnaround time.

Four characteristics of agentic AI

1. **Relies on memory:** Agentic AI can access historical data and track previous interactions.
2. **Can use tools:** It can integrate with external systems through application protocol interfaces, EHRs, and scheduling tools.
3. **Is able to plan**: It can break down workflows, plan next steps, and predict outcomes.
4. **Can be autonomous:** It can leverage memory tools and planning to operate independently.

Agentic AI is often compared to a capable executive assistant rather than a calculator or a smart word processor. Traditional AI waits for explicit instructions, much like an assistant waiting for a specific task. Generative AI is closer to an assistant who can draft emails or summarize reports once asked but still needs direction throughout the process. Agentic AI behaves more like an assistant who understands priorities and operates independently. It notices a problem, determines what needs to be done, sequences the work, and acts across systems.

For health plan leaders, the distinction matters because agentic AI begins to behave less like a tool and more like a digital operator — one that could streamline complex, cross-functional processes such as utilization management, care coordination, or provider outreach. Agentic AI is transformative not because it automates a single task, but because it enables plans to operate with greater speed, coordination, and resilience at enterprise scale.

Why agentic AI can be transformative

1. **Can ingest complexity:** Agentic AI can process multiple types of intake channels simultaneously without the need for perfect data and can thrive in messy ecosystems.
2. **Understands healthcare context:** Instead of treating everything as generic content, agentic AI understands healthcare-specific terms, regulatory requirements and clinical pathways.
3. **Collaborates with transparency:** Agentic AI supports compliance with regulatory and governance requirements by facilitating traceability, explainability, and auditability.
4. **Adapts to environmental changes:** Agentic AI systems can be quickly updated to reflect new requirements, preventing costly retraining and hard-coded failures.
5. **Scales across the enterprise:** Scalability exists across use cases and business functions, creating a cycle of transformation.

What AI adoption looks like at health plans

AI adoption has become mainstream across the health insurance industry, with 84% of U.S. health insurers reporting that they currently use AI or machine learning in some part of their operations.[1] Adoption is concentrated in high-volume, labor-intensive administrative and clinical functions, including utilization management (71% of plans), prior authorization (68%), disease management (61%), risk adjustment (57%), medical provider fraud detection (51%), claims fraud detection (50%), and sales and marketing (45%)[1]. AI deployment is largely focused on augmenting human decision-making rather than replacing it, with formal governance, bias monitoring, and human oversight remaining standard practice.

Why health plans are cautious about agentic AI

Many health plans believe AI has significant potential, but adoption is often constrained by trust rather than technology. While 84% of insurers report using artificial intelligence or machine learning, most deployments remain focused on supporting discrete functions rather than fully automating end-to-end operational processes through agentic AI[1].

Health plans generally view AI as a tool to augment human decision-making, with extensive governance, bias testing, model monitoring, and human oversight still embedded in workflows. As a result, the path to broader transformation requires plans to first build confidence through narrowly scoped, lower-risk use cases — such as document summarization, data extraction, workflow routing, or fraud flagging — where AI recommendations can be validated by humans and performance can be measured over time.

Many organizations struggle to take these steps because operational leaders remain accountable for regulatory compliance, member outcomes, and financial accuracy, making the perceived risk of errors often outweigh the potential benefits of automation. In addition, fragmented data, legacy systems, unclear governance models, and limited proof of return on investment can make it difficult to progress from isolated pilots to trusted, enterprise-scale automation.

Agentic AI use cases for health plans

For health plans, early success with agentic AI tends to come from high-volume, administrative-heavy workflows where decisions can be clearly bounded by policy, data is accessible through APIs, and performance metrics such as cycle time or denial rates are already tracked. In these areas, agentic AI tools can orchestrate multistep work under defined governance and human oversight.

A common approach is to first deploy one component: one product line, one channel, or a limited policy set. Leaders can measure performance before and after implementation and then expand the use of AI based on demonstrated results.

Use case 1: Managing prior authorizations

Health plans are developing AI solutions that operate within the framework of platforms such as Availity. These solutions can guide providers through questionnaires when they submit authorization requests and show the criteria needed to meet medical necessity for the requested procedures. Submitters can also attach electronic health records with the requests. After submission, these solutions can automatically provide an approval recommendation for each submission in less than 90 seconds when the necessary criteria are met.

Use case 2: Identifying high-risk and rising-risk members

Health plans currently use Hierarchical Condition Category (HCC) models to assess risk. Agentic AI can continuously monitor claims and EHR feeds, detect uncaptured conditions or signs of deterioration, and initiate targeted outreach to members or providers. Over time, these systems can learn which interventions are most effective at closing gaps and reducing avoidable utilization.

Use case 3: Streamlining call center operations

Healthcare contact centers handle high volumes of routine inquiries related to eligibility, claim status, prior authorization, ID cards, and similar issues. Agentic AI can manage many of these interactions through proactive notifications, conversational self-service, and agents that retrieve real-time information, verify benefits, update member records, or schedule follow-ups. Plans should define success metrics such as first-call resolution, customer satisfaction, and supervisor assist rates.

Other potential areas for agentic AI use

- Longitudinal care coordination
- Payment integrity
- Provider network management
- Claims management
- Risk and quality monitoring
- Utilization management/prior authorizations

How to successfully implement agentic AI

As agentic AI moves from experimentation to execution, health plans are moving from what is possible toward what is practical. Leaders consistently point to the following considerations as differentiators between pilots that stall and programs that scale.

Start small and scale

For health plans looking to implement agentic AI, the most effective starting point is not the most ambitious use case. Instead, start with the most practical one. Low-complexity workflows that are rules-based, repetitive, and already well understood offer the fastest path to value and the clearest return on investment. Starting small allows organizations to build confidence, refine operating models, and demonstrate impact before expanding into more complex environments like end-to-end member journeys. The discipline is in choosing the right first use case, not the most exciting one.

Decide your organization's appetite for risk

Being an early adopter of agentic AI can create competitive advantage and accelerate learning, but it also requires a willingness to manage uncertainty and invest ahead of proven benchmarks. Others may become fast followers, adopting proven patterns once guardrails and best practices are clearer. Neither approach is inherently right or wrong; what matters is aligning AI investment decisions with the organization's strategic priorities, risk tolerance, and mission.

Pair autonomy with human oversight

As agentic AI systems take on multistep tasks independently, organizations must clearly define where autonomy stops and human judgment begins. In health plan operations, not every decision should be fully automated, especially when outcomes affect clinical care, member access, or financial liability. Organizations that adopt agentic AI solutions must identify and implement human-in-the-loop workflows to ensure human oversight while maintaining efficiency. While agentic AI can significantly streamline payment integrity, prior authorizations, and call center operations, expert review remains critical for high-impact decisions and edge cases.

Build governance that scales with technology

Organizations that successfully scale agentic AI tend to treat governance as an enabler rather than a constraint. A structured review process and ongoing monitoring help address issues like model drift, bias, data misuse, and unintended consequences as systems evolve. Many leaders are establishing cross-functional AI review boards to oversee use case approval, model changes, and performance monitoring alongside traditional concerns like cybersecurity, data sharing, and liability.

Conclusion

Looking ahead, agentic AI has potential to both save costs and bolster revenue, which will free scarce resources to reimagine healthcare delivery, improve outcomes, and hyper-personalize member experience. Plans that invest in the right foundations — clear use cases, human oversight, and scalable governance — will be best positioned to turn agentic AI into a sustained operational advantage as the technology matures.

Want to see these use cases in real life?

Optum consulting partners with payers and providers to adopt and scale agentic AI across the enterprise.

Get in touch at **advisory.com/optum-support**

Related content

EXPERT INSIGHT

Payer-provider AI race series, part 1: The race explained

Read now

EXPERT INSIGHT

Payer-provider AI race series, part 2: Where the race is headed

Read now

CASE STUDY

How BCBSM used AI to improve member outreach and engagement

Read now

Endnote

1. National Association of Insurance Commissioners. NAIC AI Health Survey Report. National Association of Insurance Commissioners. Accessed July 22, 2026.

Advisory Board

655 New York Avenue NW, Washington DC 20001 | **advisory.com**