Agentic AI and the Next Frontier of Healthcare Transformation

2026

Introduction

The Health Management Academy (THMA) and Microsoft Partnership

The Health Management Academy (THMA) and Microsoft have partnered to advance health systems' AI transformation journeys through a joint **AI Transformation Research and AI Collaborative** initiative. This collaboration continues to demonstrate that Microsoft is a trusted partner for health system executives navigating the operational and strategic implications of AI adoption. Together, THMA and Microsoft are conducting co-authored research on **AI maturity**, designing and hosting **executive workshops** on real-world AI use cases across C-suite forums, and led an **exclusive AI Collaborative program** for CMIOs, CNIOs, and analytics leaders. These combined efforts aim to accelerate industry-wide learning, strengthen cross-system collaboration, and provide practical frameworks to help health systems move from AI experimentation to enterprise-scale impact

About The Health Management Academy

Since 1998, The Health Management Academy has cultivated the premier community of influential changemakers in healthcare. Our members are aligned around a common goal of improving health for all, and a core belief that partnership will accelerate progress. Our member community includes Leading Health Systems – the approximately 150 innovative integrated delivery systems with over $2B in total operating revenue – and innovative Industry Partners that are working alongside health systems to drive health forward. We power our members by building our community and fostering connections through executive peer learning. We support professional growth through talent and development. We accelerate understanding by delivering timely and actionable data and insights on key challenges. And we catalyze transformation by building alliances in areas where the power of the collective is greater than the power of one.

Agentic AI and the Next Frontier of Healthcare Transformation

This report examines how health systems are advancing along the agentic AI maturity curve, from initial adoption to enterprise-wide optimization. Each section centers on a core dimension for AI-driven transformation in health systems, including:

1. **A New Era of Healthcare AI:** Defines agentic AI and its progression toward autonomy. Explores how systems are integrating digital colleagues into workflows, aligning around outcomes, and preparing for human-agent team management.
2. **AI Value Realization: Market Goals and Constraints:** Examines the state of the market today, and how current macro pressures shape AI investment priorities. Assesses the practical barriers in governance and ROI actualization impeding scalability of AI solutions across health system enterprises and highlights key areas of opportunity to inflect transformation efforts moving forward.
3. **AI Transformation in Practice and Over the Next 3–5 Years:** Identifies where investment concentrates today and how external partners are enabling transformation amid internal capability gaps. Explores AI's evolution from experimental tool to foundational capability shaping financial performance, clinical care, and workforce models. Emphasizes near-term impact with risks amplified by scaling.
4. **Future-Proofing Your System for Agentic AI:** Outlines three interdependent imperatives for operationalizing agentic AI: embedding governance as a continuous system, building a mature data infrastructure, and developing workforce capabilities to manage AI agents.
5. **Executive AI Perspectives:** Examines the unique perspectives of AI leaders from a set of 5 diverse health systems. Highlights the current state of AI strategy and maturity within each health system, as well as their key barriers to AI transformation, and their long-term aspirations.

Methodology Overview

The insights in this report are grounded in a mixed-method research approach conducted jointly by **The Health Management Academy (THMA)** and **Microsoft** as part of the **AI Transformation Research and AI Collaborative**. The methodology combined **quantitative survey data** with **qualitative executive insights** to capture a comprehensive view of how health systems are advancing along the AI maturity curve.

The **quantitative component** drew from a national survey fielded by THMA to executives across **Leading Health Systems (LHS)**, large integrated delivery networks representing the forefront of digital and operational transformation. The survey included **30 respondents (N=30)** spanning **C-suite, clinical, IT, and analytics leaders**, from 30 unique LHS, ensuring diverse perspectives across governance, technology, and care delivery domains. The survey measured organizational maturity across several dimensions, including **strategy, infrastructure, governance, workforce readiness, and early agentic AI adoption**.

The **qualitative component** leveraged discussions and case studies from THMA's **AI Collaborative**, held in partnership with Microsoft. This multi-day forum convened **AI, clinical, and operational leaders** from LHS organizations to explore emerging practices, challenges, and frameworks for enterprise-scale implementation. Insights from these conversations were integrated with survey findings to contextualize quantitative trends and surface practical examples of innovation, governance models, and leadership strategies.

Together, these data sources provide a holistic view of how leading health systems are moving from experimentation to operationalization, defining what AI maturity looks like in healthcare and identifying the enablers of sustainable and responsible transformation.

Section 1:

A New Era of Healthcare AI

This section examines how healthcare organizations are evolving toward the agentic era by defining agentic AI, tracing its progression from assistance to autonomy, and applying the Frontier model to care delivery. It explores how health systems can embed digital colleagues into workflows, redesign structures around outcomes, and strengthen workforce readiness to manage human-agent teams, positioning leaders to drive smarter, more sustainable transformation.

Agentic AI: Healthcare's Shift from Predictive Tools to Digital Colleagues

Defining Agentic AI

AI is entering a new phase defined by **agentic systems** capable of reasoning, planning, and executing tasks with varying levels of autonomy under human oversight.

These "agents" move beyond static tools or copilots. They act as **digital colleagues**, systems that can complete end-to-end workflows, learn from feedback, and report progress back to humans. As Microsoft describes it, *"intelligence on tap"* will reshape how organizations work by blending **machine intelligence with human judgment** to create structures that are *AI-operated but human-led*.

In Practical Terms, Agentic AI Allows Health Systems to:

- Automate documentation, scheduling, and other repetitive tasks
- Deploy intelligent agents that triage patient messages or manage supply chain logistics
- Free human teams to focus on strategy, empathy, and problem-solving

---

Microsoft frames this transformation as a three-phase journey that mirrors the evolution many health systems are beginning to experience:

Phase 1	**Human with Assistant**	AI copilots help clinicians and staff complete routine work faster, like documentation summaries or inbox management.
Phase 2	**Human-Agent Teams**	Agents join care and operations teams as digital colleagues, taking on specialized tasks such as pre-visit preparation, claims review, or staffing optimization.
Phase 3	**Human-Led, Agent-Operated**	Humans set direction and priorities while agents execute business processes end-to-end, checking in only for oversight or exceptions.

Scaling Smart: How Agentic AI Rewrites the Rules of Work

Agentic AI is not simply about efficiency. It represents a new organizational model built on **hybrid human-AI teams**.

- **For patient care:** Agents can triage clinical messages, surface relevant data from EHRs, and prepare documentation, giving nurses and physicians back time for bedside care.
- **For workforce design:** Every employee becomes an *agent boss*, responsible for designing, training, and supervising digital colleagues. Leaders at the Frontier (those already operating this way) are twice as likely to say their companies are thriving.
- **For financial sustainability:** Frontier organizations report both higher productivity and higher employee satisfaction. AI-enabled processes expand capacity without adding headcount, an essential lever in an era of tight margins.

> **Bottom Line:** Healthcare's next productivity leap will come from **human-led, agent-operated systems**. The organizations that build AI literacy today will be best positioned to scale clinical capacity, reduce burnout, and sustain growth.

The Rise of Agentic AI: From Assistance to Autonomy

AI in healthcare has evolved from narrow predictive models to **generative assistants** and now to **agentic systems**, digital colleagues capable of reasoning, planning, and acting with human oversight.

Early predictive tools offered insight but not action. They identified risks, surfaced data patterns, and informed decisions, but execution remained fully human. The emergence of generative AI marked a turning point in automating documentation, scheduling, and communication. Today, agentic systems have the potential to extend beyond assistance to operate as partners within workflows, reshaping how healthcare teams function every day.

Phases of Evolution in Healthcare AI

1

AI as an Assistant

The earliest stage of healthcare AI focused on **automation**. Generative tools accelerated administrative work such as summarizing visit notes, scheduling follow-ups, and drafting patient messages. These tools saved time but largely replicated existing tasks rather than reimagining them.

2

Human-Agent Teams

The next phase introduces **digital colleagues** that share work with humans rather than simply supporting them. These agents perform discrete, well-defined tasks under clinician or manager oversight. Examples include handling prior authorizations, routing messages, or preparing discharge instructions. The result is a more balanced partnership that frees staff to focus on complex, high-value care.

3

Human-Led, Agent-Operated Workflows

In the emerging third phase, **humans set goals and agents execute**. Agents manage end-to-end workflows such as claims reconciliation, supply chain operations, or population health outreach. Humans intervene when judgment or empathy is required but no longer manage each step directly. This model represents a fundamental shift in how healthcare work is designed and led.

These phases often coexist. Many health systems are already piloting agents in some departments while relying on copilots in others. The journey is not linear but adaptive, evolving as organizations learn how to safely and effectively scale agentic systems.

A Gradual but Transformative Shift

Like the industrial and internet revolutions before it, this transformation will unfold **gradually but decisively**. Each stage requires new forms of coordination between technology and human skill.

- **Culturally,** teams must learn to trust and guide digital collaborators.
- **Technically,** organizations must integrate AI into secure and reliable workflows.
- **Operationally,** leaders must redesign oversight and accountability structures for hybrid teams of humans and agents.

These shifts take time, but they also unlock compounding gains in efficiency, accuracy, and workforce well-being.

The Frontier Model Applied to Healthcare

Defining What it Means to be Frontier

In the *Work Trend Index 2025*, Microsoft describes organizations at the "**Frontier**" as those integrating *intelligence on tap*, **human-agent teams**, and **digital labor at scale**. These firms are not merely adopting AI tools; they are redesigning the way work gets done.

Organizations at the Frontier blend **machine intelligence with human judgment**, creating hybrid structures where humans set direction, and AI systems execute repeatable, data-driven processes. The result is an organization that is more agile, scalable, and capable of delivering value faster.

For healthcare, this model offers a powerful framework for restructuring around **outcomes and workflows** rather than traditional functions or titles.

Blueprint for the Agentic Health System

Three shifts healthcare leaders should be aware of as Frontier models take shape

Restructuring Service Lines

Agentic AI creates the potential to reimagine service lines by integrating digital labor across functions such as revenue cycle, scheduling, or tumor boards. For example, AI agents can manage prior authorizations, track claim denials, or coordinate multidisciplinary meetings—allowing clinical and administrative leaders to focus on strategy and patient outcomes rather than repetitive execution.

From Org Charts to Work Charts

The *Work Trend Index 2025* predicts that organizations at the Frontier will replace traditional organizational charts with "work charts," structures organized around processes and outcomes rather than departments. For health systems, this shift means teams will form dynamically around specific goals such as reducing readmissions or improving throughput. Agents will become integral members of these teams, performing research, analysis, and administrative coordination. Success metrics will move away from siloed departmental KPIs toward shared, cross-functional outcomes. This model reflects the collaborative nature of modern clinical care, where multidisciplinary teams already align around the patient journey. Agentic AI expands that approach across the entire enterprise, creating more flexible, outcome-driven organizations.

Implications for Medical Education and Workforce Development

As AI becomes a routine part of healthcare operations, AI literacy will emerge as a core competency for clinicians, administrators, and leaders. Future healthcare education will need to extend beyond technical skills to include how to direct, evaluate, and govern AI systems effectively. Medical schools, residency programs, and leadership fellowships will need to teach clinicians how to partner with AI in decision-making, equip leaders to manage hybrid teams that include digital agents, and reinforce ethical and safety standards for human-led, agent-operated environments. This transformation resembles previous technological shifts such as the adoption of electronic health records but requires a deeper understanding of reasoning, transparency, and trust between humans and machines.

Section 2:

AI Value Realization: Market Goals and Constraints

This section examines the state of the market today, and how current macro pressures shape AI investment priorities. It assesses the practical barriers in governance and ROI actualization impeding scalability of AI solutions across health system enterprises and highlights key areas of opportunity to inflect transformation efforts moving forward.

Reality of the Market Today and It's Impact on AI Valuation

Persistent cost rigidity and continued margin compression are raising expectations for AI to deliver clear, near-term financial returns. At the same time, persistent workforce shortages and clinician burnout are sharpening scrutiny around whether AI meaningfully reduces administrative burden and protects clinical capacity.

In practice, these pressures, as well as governance and measurement constraints, are driving a pragmatic AI strategy, with leaders emphasizing near-term financial and operational returns while longer-term clinical ambitions remain aspirational.

How executives rank the following priorities when assessing AI investments today, *by level of importance*

Percentage of respondents

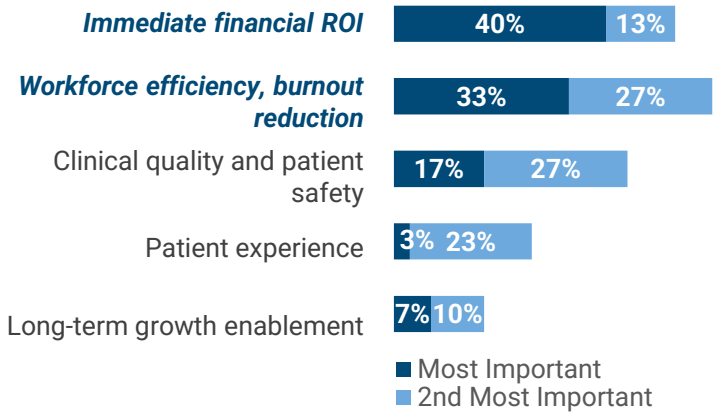

> **"We know certain market factors will continue to hurt our P&L. The promise of cost savings from AI is essential."**
>
> - Advanced Analytics Leader, LHS

Which operational domains have shown the most quantifiable ROI from AI investments in your organization?

Percentage of respondents

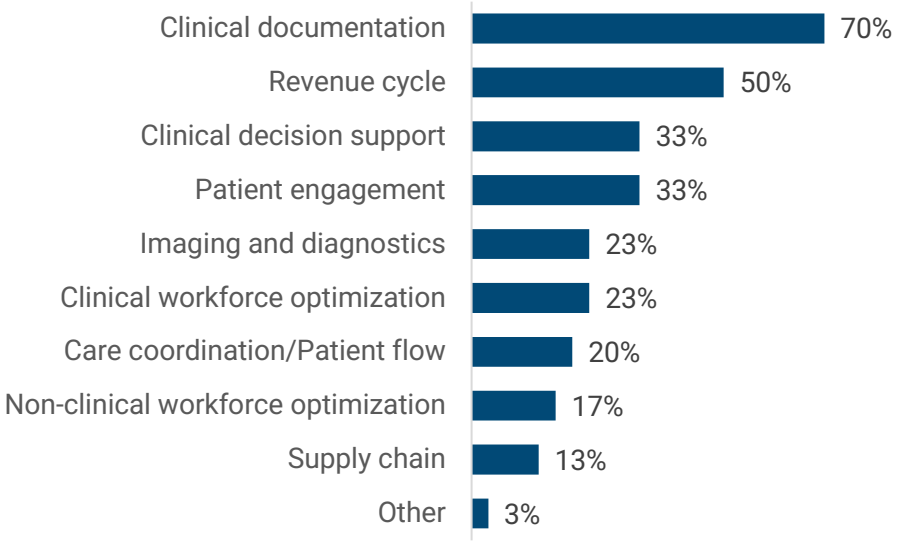

Health systems most frequently report quantifiable AI impact in administratively dense, high-volume workflows such as clinical documentation (70%) and revenue cycle management (50%). This concentration reflects where AI is delivering value today, but also how organizations define and prioritize "impact" under current financial and operational pressures.

Importantly, this pattern does not suggest that meaningful impact is absent in other domains. Rather, it highlights the alignment of these administrative use cases with enterprise strategic priorities, such as revenue lift and physician wellbeing.

This also illustrates the relative ease of measuring outcomes at scale in these domains, given their enterprise commitment. In contrast, more clinically nuanced applications often face greater challenges in generating standardized, enterprise-wide evidence of impact, contributing to their lower visibility despite ongoing investment and localized value creation.

Net, Financial ROI Actualization Still An Aspiration For Most Domains

Comparison of actively scaling AI use cases, their financial ROI, and budget limitations

Percentage of respondents indicating "yes" for each attribute (i.e., scaling, direct/financial ROI achieved, and budget constraints experienced)

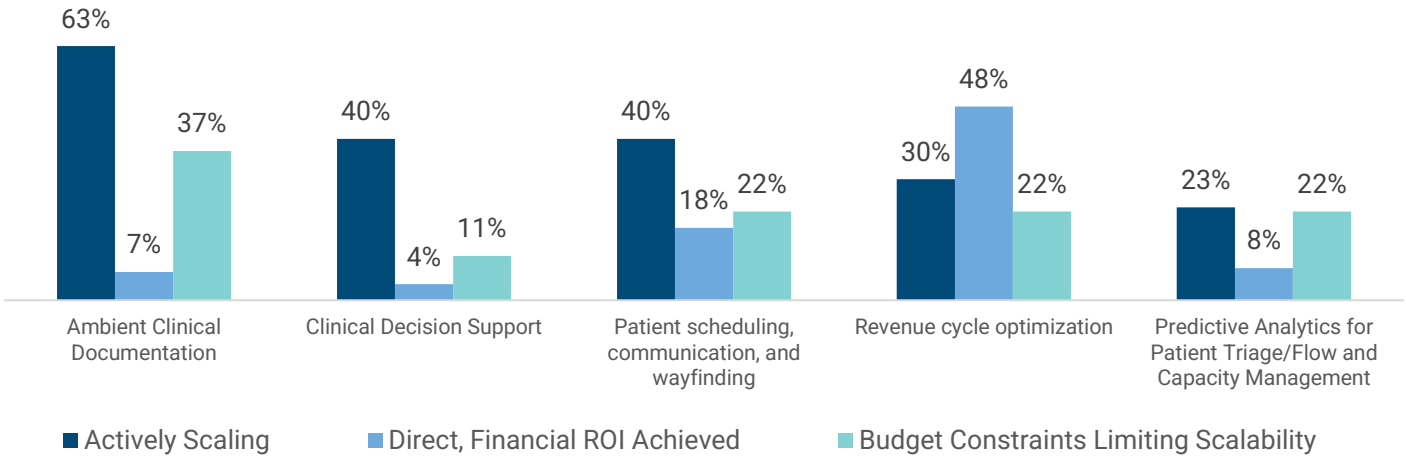

Despite clear momentum behind the most widely adopted AI applications, barriers persist in translating scale into validated, positive financial ROI. Even among the most frequently cited use cases in today's market, including ambient clinical documentation, clinical decision support, and patient scheduling or communication, health systems report limited net financial return alongside persistent budget constraints. This reflects both the genuine difficulty of realizing near-term financial gains from many AI investments and a narrower framing of what constitutes financial ROI in practice.

In this context, executive commentary suggests emerging recognition that traditional revenue-based measures may understate AI's financial contribution, with increasing attention to alternative forms of cost avoidance, such as liability reduction, and risk mitigation. While not yet systematically measured or widely validated, these perspectives point to an evolving interpretation of financial ROI, particularly as organizations assess more advanced and agentic AI capabilities.

With limited revenue gains, executives are reframing ROI around **cost avoidance** and **risk mitigation.**

"We knew in our gut there was ROI. Engaged an actuary – validated a reduction in liability. An actuary should come in every year and validate for Finance." - CIO, LHS

Divergent Views on AI's Enterprise Role Complicate Governance at Scale

As AI expands into more sensitive and role-specific tasks, health systems are progressing along different maturity paths in how AI is operationalized and governed. Across the market, executives describe AI's current impact as ad hoc (7%), operational (30%), systemic (50%), or transformational (10%), reflecting substantial variation in enterprise readiness and strategic posture. The limited share of organizations classifying AI as transformational suggests that, for most health systems, AI is still being deployed as a set of targeted capabilities rather than as a core enterprise platform. Notably, the variation in use case–driven deployments can inhibit the development of consistent governance practices and slow progress toward scalable, enterprise-wide implementation.

Given your organization's current operationalization of AI, how would you classify the impact of AI across the enterprise?

Percentage of respondents

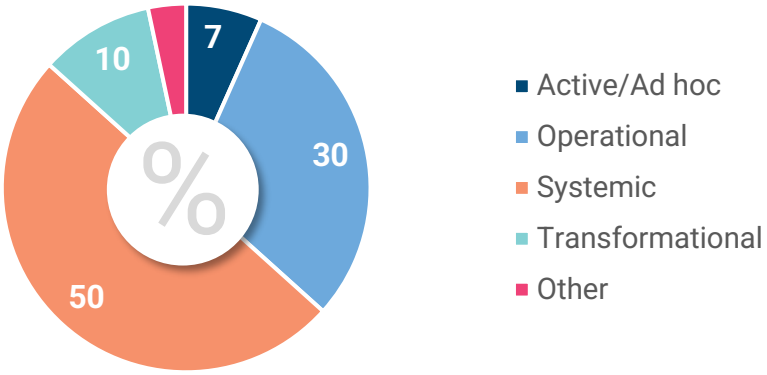

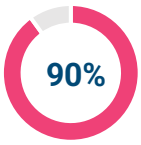

90% of executives classified their organizations' AI operationalization as something other than core to enterprise transformation

Has your organization established any of the following solutions for monitoring AI-model performance and impact at scale?

Percentage of respondents

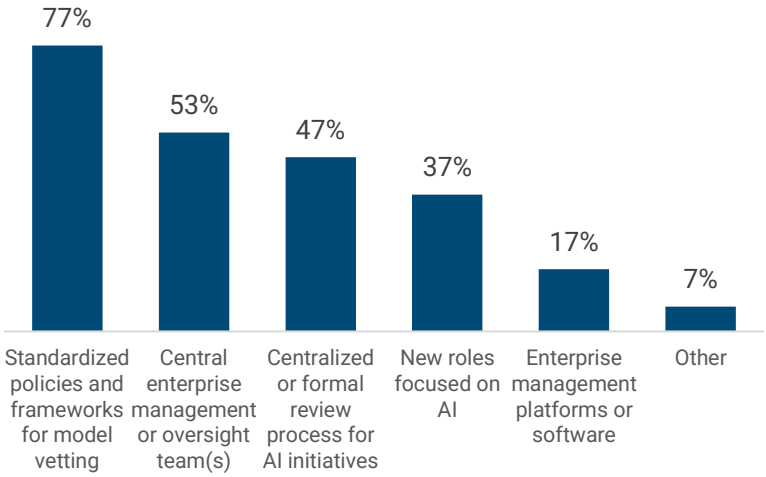

Consistent with these maturity differences, health systems report greater adoption of governance mechanisms focused on AI intake and approval than on ongoing performance management at scale. While a majority have implemented standardized policies for model vetting (77%), centralized oversight teams (53%), and formal review processes (47%), far fewer report deploying enterprise management platforms (17%) or other tooling designed to support continuous monitoring and accountability across AI portfolios. This pattern suggests that, across the market, governance capabilities are generally more developed at the front end of the AI lifecycle than at the point of sustained operation and measurement. As AI ecosystems grow in size and complexity, these gaps increasingly constrain the ability to evaluate performance, manage risk, and sustain confidence in AI-enabled workflows over time.

What are your top challenges when trying to measure AI model performance and outcomes at scale?

Percentage of respondents

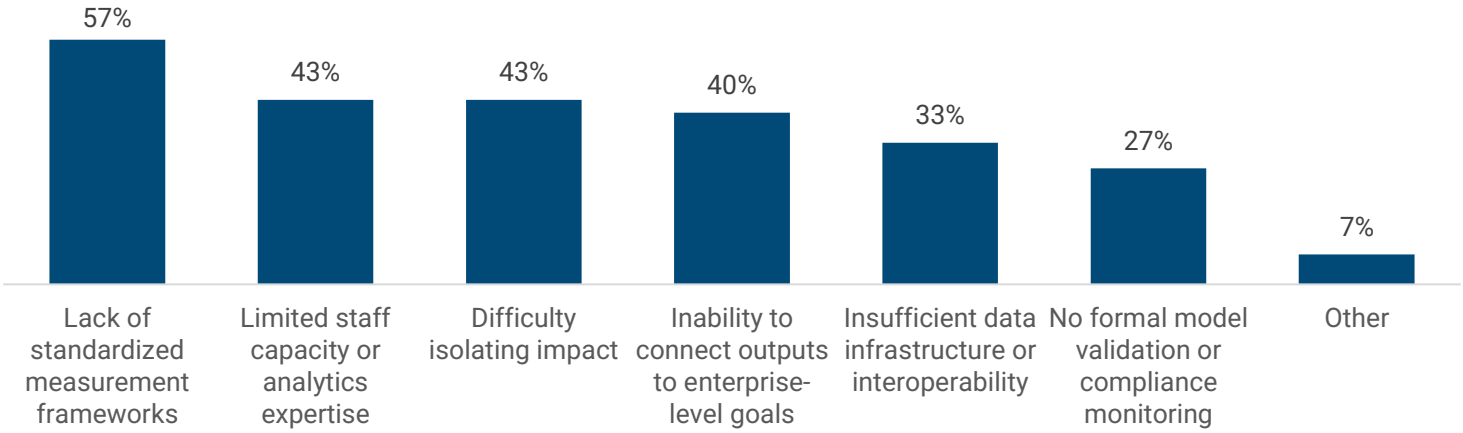

The most frequently cited challenge in measuring AI performance at scale is the lack of standardized measurement frameworks (57%), followed by difficulty isolating AI's impact from concurrent initiatives (43%) and connecting AI outputs to enterprise-level goals (40%). These findings reinforce that, as AI becomes more embedded across clinical and operational workflows, attributing discrete outcomes becomes increasingly complex. Measurement challenges intensify not because governance structures are absent, but because existing approaches were not designed to support continuous evaluation across a growing portfolio of AI use cases.

Consistent with these results, CIOs and other digital and AI leaders participating in the AI Collaborative repeatedly identified the lack of scalable, enterprise-level mechanisms for ongoing performance evaluation as a persistent market gap, particularly as AI initiatives move beyond pilot deployments and proliferate across the organization.

What are your top challenges in managing the growing ecosystem of AI models and agents across your organization?

Percentage of respondents

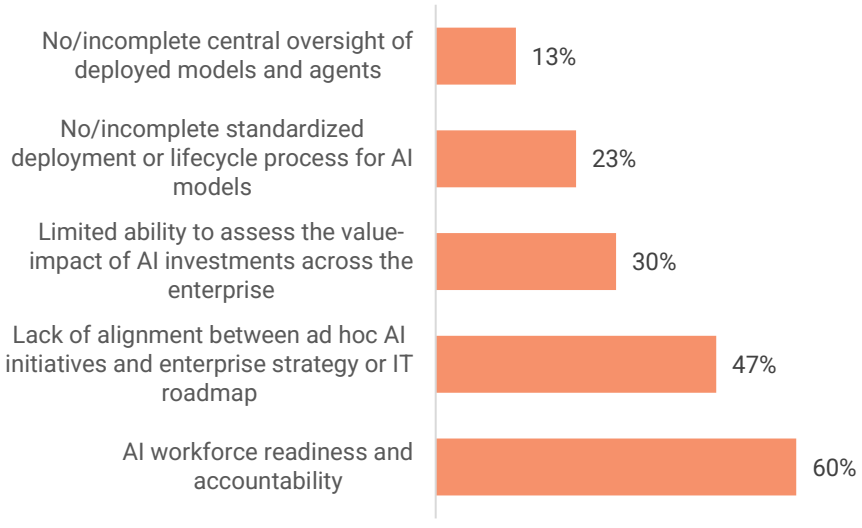

As AI portfolios expand, leaders report growing challenges related to accountability, workforce readiness, and sustaining alignment with enterprise priorities. Sixty percent of respondents cite AI workforce readiness and accountability as the primary challenge in managing expanding AI ecosystems, while nearly half (47%) point to difficulty maintaining alignment between individual AI initiatives and enterprise roadmaps. These findings suggest that governance constraints increasingly emerge downstream, as AI moves beyond pilot deployments and into broader operational use, requiring new operating models that extend responsibility for AI oversight beyond IT alone.

Workforce Readiness and Data Architecture: Core Enabler

Views from LHS AI leaders on the importance and impact of AI readiness in the workforce

Percentage of respondents indicating "yes" for each attribute/question

60% of executives **believe building an AI-enabled workforce (e.g., skilling/reskilling)** is one of the top critical future growth opportunities needed to enable AI transformation (top ranked opportunity)

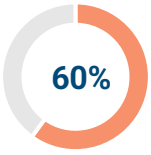

Yet 60% of executives **believe capacity for skilling/reskilling for AI (as well as unclear roles for accountability outside of IT)** is a top challenge in managing the growing ecosystem of AI models/agents across the organization (the most frequently cited challenge).

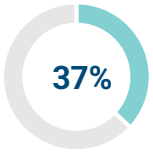

Only 37% of executives stated their organizations are **reorganizing business units or functions to support AI-enabled workflows (including AI skilling/reskilling)** in response to, and to enable, enterprise-wide AI adoption.

Across both the survey data and executive discussions, workforce readiness consistently emerges as the central constraint shaping AI adoption and scale. Sixty percent of health system leaders cite building an AI-enabled workforce, including skilling and reskilling, as a top future growth opportunity, while an equal share report limited organizational capacity and unclear accountability for AI outside of IT. Together, these findings indicate that gaps in skills, roles, and ownership continue to limit the ability to operationalize governance structures and sustain AI deployment beyond isolated use cases.

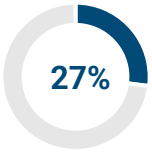

Only 27% of executives stated their organization is **redesigning core systems and data architecture to support AI** in response to, and to enable, enterprise-wide AI adoption.

In parallel, data architecture remains a material constraint on AI scalability across the market. Only 27 percent of health systems report actively redesigning core systems and data architecture to support AI at scale, despite growing recognition of the need for governed data platforms that enable analytics, model training, and auditable data sharing. This gap reflects the structural and resourcing challenges associated with modernizing legacy data environments, rather than a lack of awareness of their importance.

Taken together, workforce readiness and data architecture represent interdependent foundational capabilities. Progress in either domain is necessary but insufficient on its own, as limitations in skills, accountability, or data infrastructure can constrain the ability to deploy AI reliably and transparently at scale

"Training on efficient and appropriate use must take into account variability of readiness and previous exposure."

- VP of Operations, LHS

"Try to assure the foundation for effective AI is in place. This means making sure your data is thoughtfully structured, priority use cases are identified, and governance structures can scale."
- Analytics and AI Leader, LHS

Section 3:

AI Transformation in Practice and Over the Next 3–5 Years

This section examines where leaders are concentrating their AI investment and execution efforts today. Workforce skilling, predictive analytics, and patient-facing automation emerge as the most critical near-term growth opportunities, supported by new enterprise governance and operating model changes. At the same time, AI investment decisions are being driven by immediate financial ROI and workforce efficiency gains rather than long-term growth or patient experience, reinforcing a pragmatic, operations-first mindset. Most organizations expect external technology partners to play a central role in enabling AI transformation, driven by internal capability gaps, the need for speed and scale, and the importance of integration, security, and alignment with major ecosystem platforms.

Priority Areas of AI-Enabled Transformation

Across health systems, AI-enabled workforce skilling and reskilling clearly stands apart as the most consistently cited next growth opportunity, identified by 60 percent of executives. This prominence reinforces the view that foundational workforce capabilities are a prerequisite for advancing more complex AI-enabled transformation efforts. Beyond workforce readiness, a broad set of use cases cluster in the mid-range, including predictive analytics for care management (37%), patient-facing conversational AI (37%), and agentic AI models supporting clinical and operational workflows (33%), suggesting distributed interest rather than a singular near-term transformational focus.

Notably, more clinically intensive and infrastructure-dependent applications, such as imaging and diagnostics, clinical decision support, and ambient monitoring (beyond documentation), are cited less frequently, while AI for treatment discovery and R&D remains a low near-term priority. The distribution reflects a deliberate sequencing of AI investment, with organizations prioritizing workforce and operational readiness before scaling more advanced, autonomous, or clinically embedded applications.

What LHS AI leaders believe will be the next critical growth opportunities for AI transformation at their health system

Respondents selected their top 3 growth opportunities

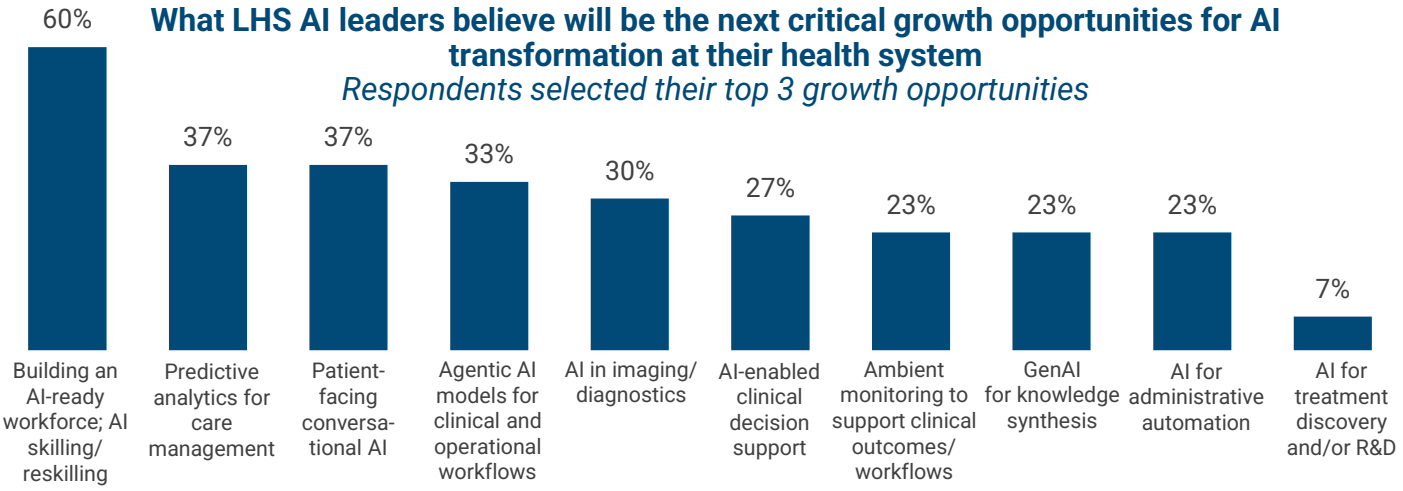

Organizational Shifts Supporting AI Adoption

To support these evolving priorities, health systems report a range of organizational and governance shifts. The most common response is the establishment of enterprise-wide AI governance or dedicated transformation leadership, cited by 70 percent of respondents, followed by formal integration of AI into enterprise strategy and investment planning (47%). Fewer organizations report deeper structural changes, such as reorganizing business units to support AI-enabled workflows, integrating non-physician clinical leaders into decision-making, or redesigning core systems and data architecture. These findings suggest that while strategic intent and oversight structures are advancing, broader organizational and infrastructural changes required to support AI at scale remain uneven.

Which of the following shifts has your organization made (or is planning) in response to, and to enable, enterprise-wide AI adoption?

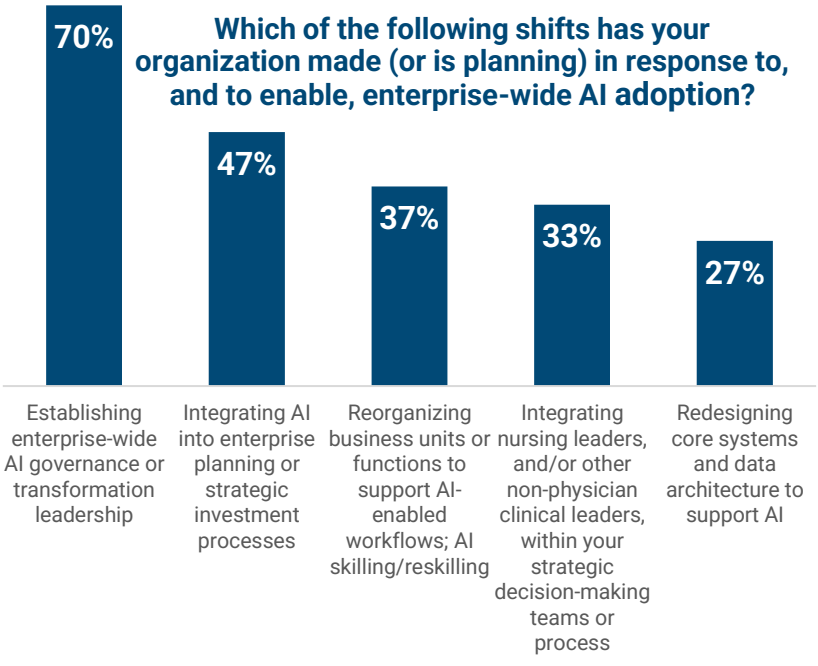

Agentic AI Still Nascent, Yet Gaining Strategic Momentum

Extent of Health System Adoption of Agentic AI Tools
Percentage of respondents

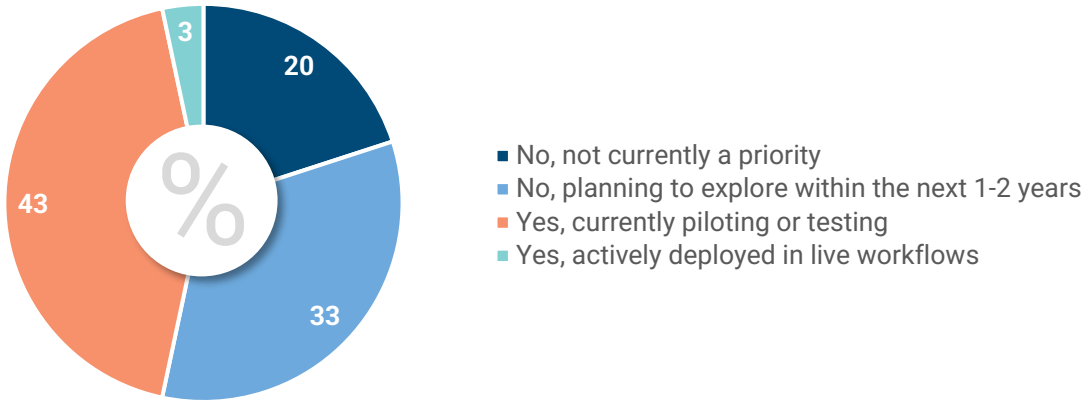

Despite growing interest, adoption of agentic AI across health systems remains early-stage. While 43 percent of respondents report piloting or testing agentic AI tools, only 3 percent indicate active deployment in live workflows, underscoring the gap between experimentation and operationalization. An additional 33 percent report no current plans to explore agentic AI within the next one to two years, suggesting that for many organizations, these capabilities remain contingent on more foundational readiness. Taken together, the data indicate that agentic AI is gaining strategic attention but has not yet transitioned into a broadly scalable or embedded operational capability across the market.

While near-term adoption remains limited, health system leaders express strong conviction in the longer-term impact of agentic AI. Sixty percent of respondents agree or strongly agree that agentic AI will meaningfully disrupt or improve the provider–patient experience, with similar optimism around gains in back-end productivity (57%). Nearly half also anticipate deeper human–AI collaboration over the next three to five years, signaling expectations that agentic models will increasingly augment, rather than replace, clinical and operational roles. This divergence between current deployment and future belief highlights a market dynamic in which agentic AI is viewed as a strategic end-state, dependent on advances in workforce readiness, governance, and data infrastructure to move from concept to sustained value.

How Health System AI Leaders View the Impact of Agentic AI Over the Next 3-5 Years
Percentage of respondents

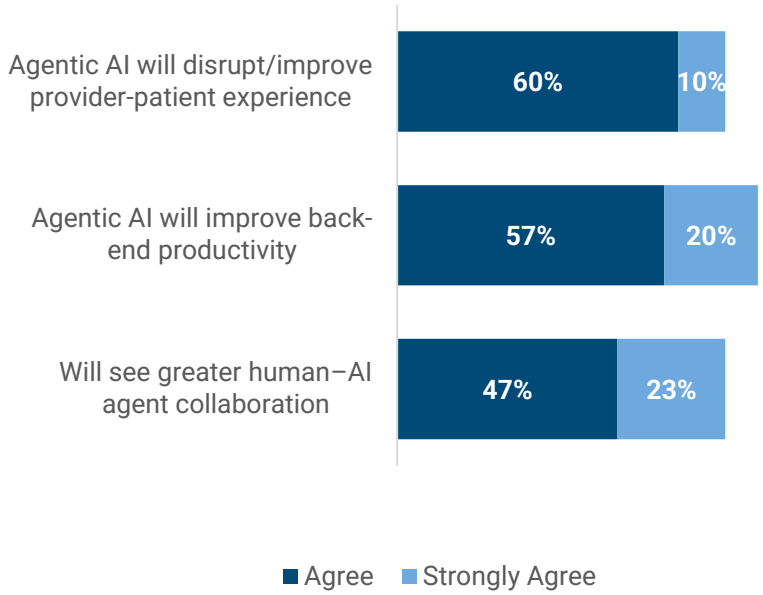

Leaders Expect AI to Become Foundational Across Enterprise Strategy and Operations

Leaders are broadly optimistic that AI will become a foundational enterprise capability over the next three to five years, with particularly strong conviction around workforce efficiency and clinical impact. However, this confidence stands in contrast to today's more cautious posture toward clinical AI deployment, where investment and execution remain limited. This disconnect between long-term belief and near-term commitment reflects persistent concerns around validation, workflow integration, and risk, echoed consistently in executive discussions.

Over the next 3–5 years, to what extent do you agree with the following statements about AI's role in transforming healthcare at your organization?

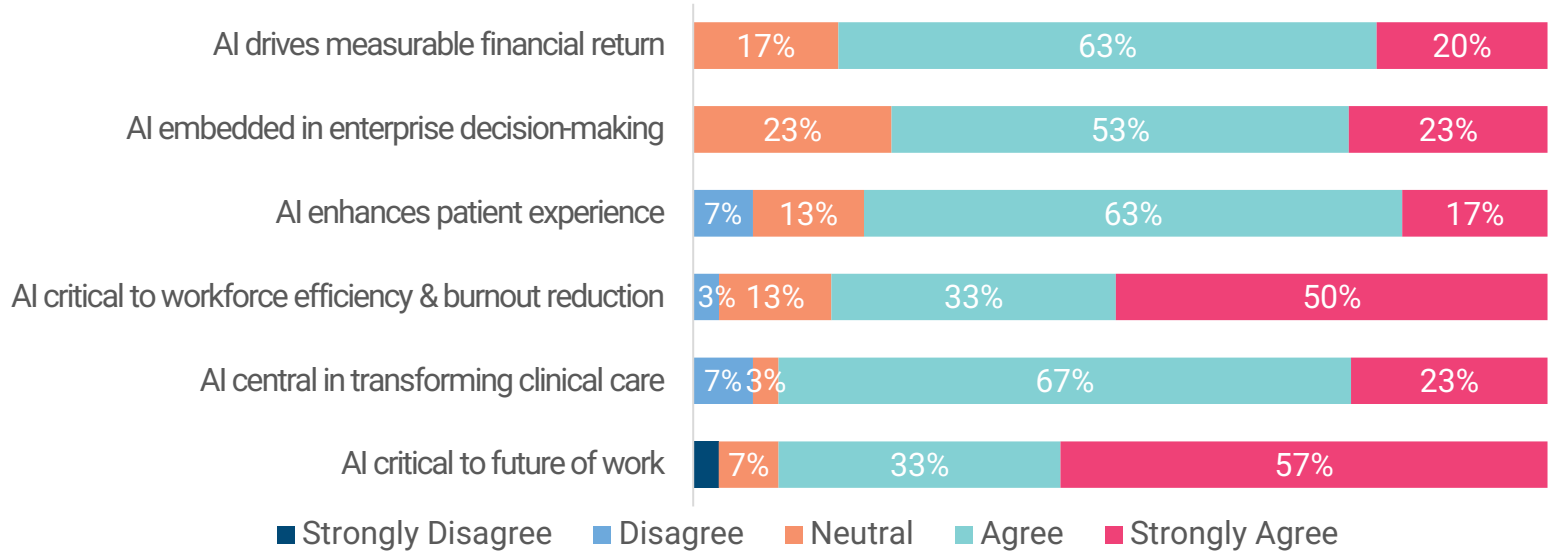

Agentic AI Expected to Reshape Collaboration, Productivity, and Patient Experience

Leaders also express strong long-term optimism about the role of agentic AI, with majorities agreeing it will meaningfully improve productivity and enable deeper human–AI collaboration over time. This optimism reflects belief in the potential of more autonomous AI models to reshape how work is performed across the enterprise. At the same time, earlier findings highlight that adoption remains early and uneven, with governance, workforce readiness, and measurement limitations continuing to temper near-term deployment. Together, the data suggest agentic AI is viewed less as an immediate solution and more as a strategic end state contingent on foundational capabilities that still need to be established across most organizations.

Over the next 3-5 years, to what extent do you agree with the following statements about Agentic AI's role in transforming business at your organization?

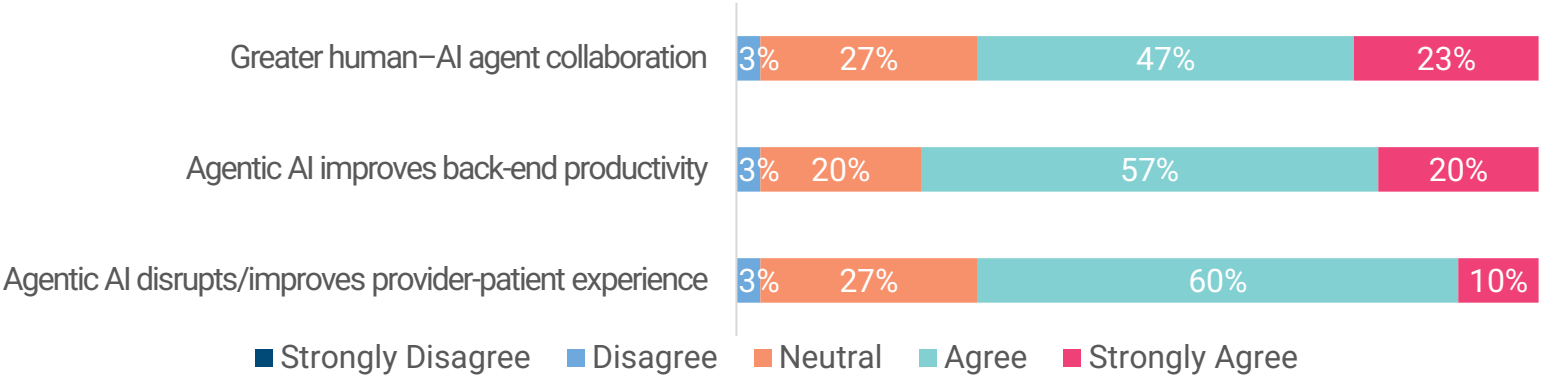

Stakeholders Expected to Experience the Greatest Impact: Operational Staff, Patients, and Clinicians

Over the next three to five years, health system leaders expect AI's most immediate and visible impact to be felt by operational support staff, patients, and frontline clinicians, reflecting early traction in high-volume, workflow-intensive domains. Nearly three-quarters of respondents (73%) anticipate the greatest impact among operational staff, where automation in areas such as scheduling and revenue cycle is already reducing manual effort and improving throughput. Patients follow closely (67%), with anticipated improvements in access, navigation, and communication, while physicians and nurses (60% each) are expected to benefit primarily through reduced administrative burden and workflow support rather than fundamental role redesign. In contrast, relatively few leaders expect AI to directly impact executive or clinical leadership roles in the near term, reinforcing a view of AI as an operational and care-delivery enabler rather than a leadership substitute.

Which stakeholder groups will experience the greatest impact from your organization's AI strategy over the next 3–5 years?

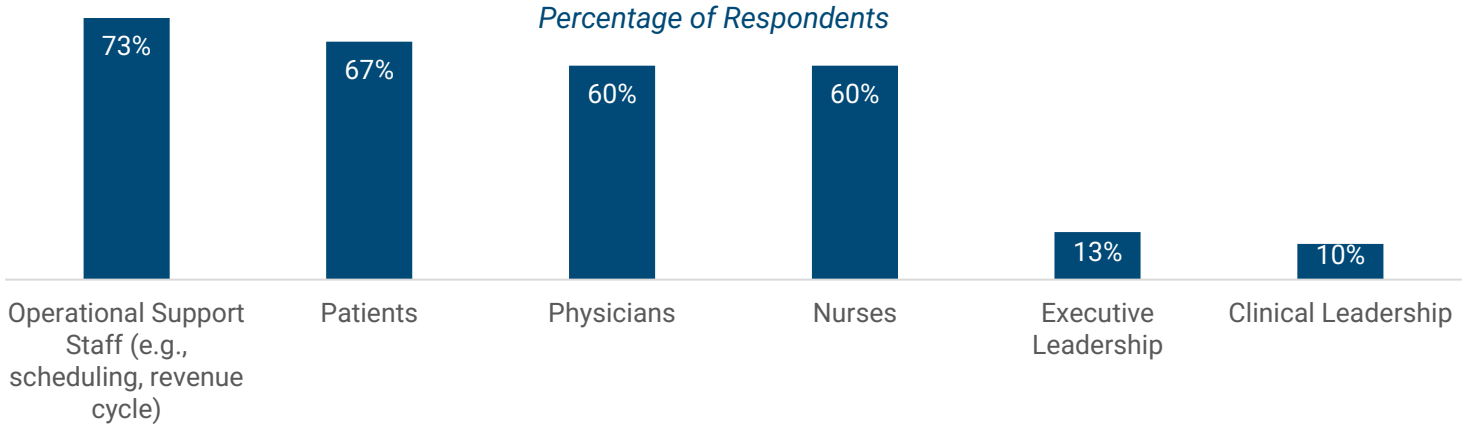

Despite these expectations, leaders underscore that realizing these impacts is contingent on managing significant risks as AI adoption scales. Patient safety and clinical accuracy and data privacy and cybersecurity are the most frequently cited concerns (60% each), reflecting heightened sensitivity as AI becomes more embedded in clinical and operational workflows. Algorithmic bias and equity considerations (40%), regulatory or legal exposure (30%), workforce disruption or resistance (23%), and the long-term financial sustainability of AI investments (23%) further reinforce that AI's value is not assumed. Instead, leaders view meaningful patient and workforce impact as dependent on sustained attention to governance, trust, and accountability alongside technological advancement.

What specific AI risks are top-of-mind for your organization?

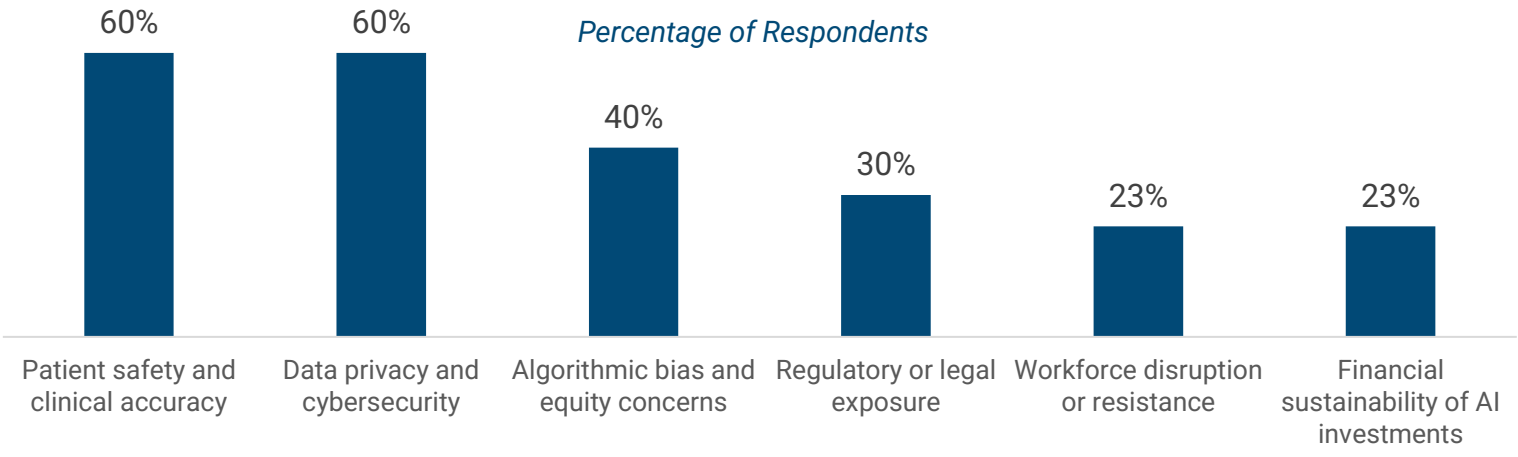

To what extent do you expect your organization's future AI transformation efforts to be supported through strategic partnerships with large technology or platform companies (e.g., Microsoft, Epic, Google, etc.)?

Percentage of Respondents

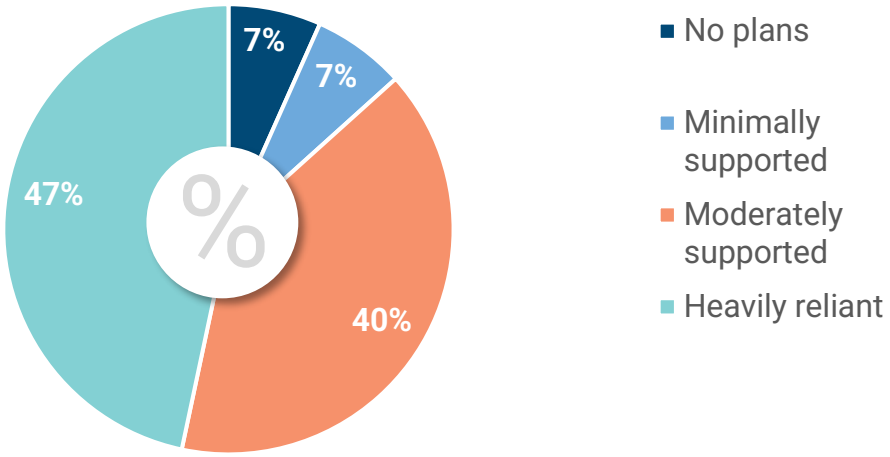

Role of External Partnerships in AI Transformation

External partnerships are expected to play a significant role in AI transformation for most organizations. Nearly half of respondents (47%) anticipate being heavily reliant on large technology partners, signaling an expectation of deep external dependency to advance AI initiatives. An additional 40% expect moderate support from external partners, reflecting a blended model in which responsibility and ownership are shared between partners and internal teams. Only a small minority report minimal reliance on partners (7%), and an equally small share indicate no plans to engage external partners at all, a rare approach that is likely tied to unique internal investment strategies, capabilities, or long-standing in-house expertise.

Drivers Behind External Partnership Decisions

Decisions to engage external partners are driven primarily by internal resource and capability constraints, including limited expertise, talent, and capacity to design, deploy, and scale AI independently. Many organizations view partners as a way to supplement internal skills while providing strategic guardrails. Speed and time-to-value are also major motivators, with leaders emphasizing the need to accelerate deployment, move more quickly from pilots to production, and achieve faster ROI through proven playbooks. Integration, interoperability, and security considerations further shape partnership decisions, as organizations seek partners that can align with existing EHRs and infrastructure while mitigating risk and ensuring compliance. Finally, strategic alignment and co-innovation remain important drivers, with leaders expressing a desire for shared vision, access to best practices, and ongoing innovation, particularly through alignment with major ecosystem players such as Epic, Microsoft, and Google.

Section 4:

Future-Proofing Your System for Agentic AI

This section examines three interdependent imperatives for operationalizing agentic AI: embedding governance as a continuous system, building mature data infrastructure, and developing workforce capabilities to manage AI agents. Organizations that address these dimensions in concert will achieve faster deployment and stronger compliance.

This section draws on insights from The Health Management Academy and Microsoft's 2025 AI Collaborative, held September 8–10, which convened AI leaders from Leading Health Systems (LHS) across the United States to discuss emerging strategies, challenges, and opportunities in enterprise AI transformation.

Governance as a Living System

Leading health systems recognize that effective **AI governance** can no longer function as an isolated, periodic checkpoint. Instead, oversight must operate as a **continuous and adaptive system** embedded into the fabric of daily operations. This requires governance to act as less of a gate and like a **guardrail**, dynamically responding to the pace of AI development rather than trailing behind it.

For healthcare, the conversation has transcended beyond use case examples and proof-of-concept demonstrations. The question is no longer whether agentic AI will transform clinical workflows, but **how organizations will manage that transformation responsibly**. To design and enable nimble operational governance frameworks, leaders must address 3 foundational dimensions:

Policy-based controls within ecosystems	Architectures that delineate responsibility	Integrated audit and provenance capabilities
Establish clear protocols for agent registration, access privileges, and real-time monitoring. Defining which agents can interact with what data, under what conditions, and with what level of human oversight enables enforcement through technical controls rather than just procedural documentation.	Create decision frameworks that map human-versus-agent responsibilities across different risk classes. High-risk clinical decisions may require human-in-the-loop validation, while lower-risk administrative tasks can operate with human-on-the-loop monitoring. Making these boundaries explicit and auditable ensures at-scale accountability.	Build systems that automatically capture agent decision-making, behavior, and intervention points to satisfy both compliance and safety standards. Embedding provenance into every transaction creates an auditable trail that supports both accountability and continuous learning.

Data Architecture and Infrastructure Maturity

Data **infrastructure maturity**, not algorithmic sophistication, will determine who succeeds with agentic AI. Consolidated, cloud-based data platforms are replacing traditionally fragmented data warehouses, permitting **unified governance** and faster model iteration. Additionally, centralized environments allow governance rules, model training, and real-time analytics to pull from a single source of truth.

Health systems must build data structures that make AI **operationally sustainable, clinically compliant, and financially viable.** Clear data frameworks and interoperability standards can reduce integration friction and promote access to AI capabilities across an organization. Fragmented ownership, inconsistent metadata, and siloed governance remain **primary barriers** to scalability and compliance. When clinical, revenue cycle, and operational data live separately with **different stewardship and quality standards**, enterprise-wide deployment of fully integrated AI models becomes nearly impossible.

The path forward requires **enterprise-level data architectures** that integrate three key capabilities:

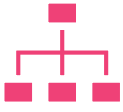

Governance by design

Build access controls, role delineation, and compliance guardrails directly into data platforms

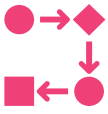

Lineage and observability

Track data provenance, behavior, and decision-making to meet regulatory requirements.

Real-time monitoring

Detect data drift, bias, and outliers as they emerge for rapid response

Workforce, Culture, and Capability Development

Agentic AI requires a **fundamentally reframed** perspective on AI. Agents should not be seen as software to deploy, but **digital colleagues** to manage. Like human workers, they require supervision, credentialing, and performance oversight defined by risk class. The challenge lies within assigning **clear ownership**: who supervises each agent, what credentials does it need for specific functions, and what level of review does each risk class require?

Addressing these questions requires capabilities most health systems have not yet acquired. Sixty percent identify **workforce upskilling** as their top challenge in managing AI ecosystems, and fewer than one-quarter have formally defined human-to-agent ratios or escalation protocols. Leading organizations are moving from continuous human-in-the-loop review toward **human-on-the-loop** models, post-event sampling, and escalation protocols. As a result, these systems maintain oversight without bottlenecking workflows. To build workforce capacity that matches the pace of AI adoption, leaders must address the following:

1. **Simulation-based training and trust calibration**
   Deploy practice-in-place programs that allow staff to interact with agents in controlled environments before deployment. Trust calibration is particularly critical; human-in-the-loop compliance can degrade within weeks as users begin rubber-stamping agent outputs. Training must help teams understand when to rely on AI outputs and when to escalate.

2. **Contextualized AI literacy**
   Develop tailored training approaches for different workforce segments. For example, clinical staff will need different training focused on reviewing agent recommendations compared to front-desk roles. Building prompt engineering skills empowers users to solve their own problems, even in non-technical roles.

3. **Institutionalized workforce management frameworks**
   Establish structures for AI workforce management, real-time performance tracking, and discipline-specific credentialing. Performance tracking should include mechanisms for feedback and reporting, enabling rapid identification of drift, errors, or workflow friction.

> **"Managing agents will require skills more akin to managing staff than managing technology."**
>
> –Chief Information Officer, LHS

Section 5:

Executive AI Perspectives

Vignettes from Diverse Health Systems

This section examines the unique perspectives of AI leaders from a set of 5 diverse health systems, representing the strategies of vertically integrated delivery networks, academic medical centers, specialty hospitals, and regional health systems. It highlights the current state of AI strategy and maturity within each health system, as well as their key barriers to AI transformation, and their long-term aspirations.

Executive Perspective: Multi-Regional Integrated Delivery Network

System and Executive Description: This system reflects a large, multi-regional integrated delivery network with a vertically integrated health plan, and a strong legacy in clinical standardization and digital enablement. AI strategy and execution are led through a centralized services organization. The executive interviewed holds responsibility for clinical informatics and works in close coordination with enterprise digital, analytics, and operational leaders.

$18b

In total operating revenue

Present in **6** states

1

Current State of AI Strategy & Maturity

AI is treated as a core, multi-year strategic priority and enterprise capability that supports clinical delivery, operational stability, and workforce sustainability across a vertically integrated system that includes both provider operations and a health plan. The system's AI strategy is grounded in the view that long-term viability depends on expanding clinical and operational capacity without proportionally expanding staffing. As a result, AI is framed less as a cost-containment tool and more as infrastructure for managing scale, complexity, and access.

The areas demonstrating the most tangible impact today are ambient clinical documentation and internal operational workflows. Ambient documentation has achieved broad physician adoption and is actively expanding into nursing, with impact measured primarily by reduced documentation burden and increased time available for patient care. This capability is viewed as foundational rather than experimental, and leadership expects it to persist as a baseline layer of clinical infrastructure.

Beyond clinical documentation, the system has demonstrated advanced maturity through deployment of agentic AI in internal operations. During a major EHR implementation, an agentic workflow was used to autonomously triage and route IT support tickets, determining priority, assigning work to appropriate teams, and requesting additional information when needed. After validating reliability, the organization allowed the system to operate autonomously for most tickets, freeing informatics staff to focus on direct clinical support. This use case represents their current focus on low-risk domains.

In parallel, the organization is actively investing in further exploration of clinically adjacent workflows, particularly care management and patient outreach, where AI can help manage information flow, triage, and coordination across provider and health plan functions. In these areas, agentic AI is being evaluated as an orchestration and prioritization tool rather than as a source of clinical decision-making. Fully autonomous clinical decisions remain explicitly out of scope.

"We look at AI as a way to manage scale and complexity, not as a way to replace people."

"The technology is moving faster than our ability to redesign workflows around it."

"The goal isn't autonomy for its own sake—it's better support for the people delivering care."

Executive Perspective: Multi-Regional Integrated Delivery Network (Continued)

2

Key Barriers to AI Transformation

The most significant technical barrier to broader AI adoption, and agentic adoption specifically, is reliability, particularly the risk of hallucination in workflows involving clinical knowledge or patient interaction. Prior experience with inconsistent AI outputs has reinforced the need for strong guardrails and conservative expansion of autonomy. From an operational perspective, workflow design, especially in ambient documentation for nursing, is also a central constraint. Nursing workflows involve real-time task tracking, handoffs, and cognitive checkpoints that are difficult to automate through ambient documentation without rethinking the underlying process.

Importantly, clinician resistance is not viewed as a primary barrier. Demand for AI-enabled tools is high, and in some cases exceeds the organization's ability to deploy solutions safely and responsibly. As a result, governance capacity, validation, and workflow readiness define the pace of progress.

3

Strategic Aspirations and Moonshots

Longer-term aspirations focus on using AI and agentic capabilities to support clinicians and care teams at scale while preserving human judgment and accountability. This includes expanding agentic orchestration in care management, integrating AI more deeply into coordination between provider and health plan functions, and continuing to build on ambient technologies to reduce administrative burden across care settings.

The organization does not articulate a vision for autonomous clinical decision-making. Instead, its longer-term direction centers on AI as an enabling layer that surfaces information, manages complexity, and supports clinicians in making better decisions, rather than making those decisions on their behalf. Progress toward this vision is expected to be incremental and domain-specific, with autonomy expanding only where risk is low and value is clearly demonstrated.

Key Takeaways

1. **AI delivers the most value when treated as core infrastructure rather than isolated pilots**. Embedding AI into essential workflows enables scale and sustainability beyond one-off efficiency gains.
2. **Agentic AI is viable today in bounded, non-clinical domains.** Starting with low-risk operational use cases allows systems to build trust and governance before exploring broader autonomy.
3. **Workflow redesign is a primary constraint.** Clinician demand for AI-enabled tools is strong; the speed of progress depends on health systems' ability to restructure workflows to absorb automation.

Executive Perspective: Large Academic-Affiliated Health System

System and Executive Description: This case study reflects a large academic medical center operating within a major public research university. The executive interviewed works within a dedicated AI and clinical decision-support team, tasked with evaluating, validating, and governing the use of AI tools in clinical and operational contexts. The role emphasizes safety, performance standards, and alignment with institutional and academic expectations, and operates in coordination with enterprise IT, clinical operations, or digital strategy functions.

$8b

In total operating revenue

Present in **1** state

1 **Current State of AI Strategy & Maturity**

AI was described as a work in progress and an active strategic priority, with leadership explicitly acknowledging that the organization is "building the bridge while crossing it." The immediate focus has been on establishing governance, safety guardrails, and monitoring mechanisms to enable responsible AI use, while simultaneously advancing a small number of practical deployments. Rather than pursuing broad AI scale-up, the strategy emphasizes learning through execution and selecting a limited set of use cases, deploying them deliberately, and refining governance and operating models based on real-world experience.

The areas where AI is currently delivering the most impact are revenue cycle operations, site-of-care decision support, and capacity management. In revenue cycle, AI is used to identify and prioritize accounts at risk for denial or delayed payment. These tools do not submit claims or take action independently; they surface specific accounts and risk indicators so staff can intervene earlier and more consistently.

In site-of-care and capacity planning, AI models are used to analyze patient characteristics, historical utilization, and operational constraints to inform decisions about inpatient versus outpatient care, bed allocation, and procedural scheduling. The models provide recommendations or risk stratification, but final disposition decisions remain with clinicians and operational leaders. The value of these applications lies in reducing variability and improving consistency, not in replacing human judgment.

Agentic AI testing is only prominent now in back-end functions and scheduling, human-in-the-loop oversight is maintained by design, not as an interim step. Yet the system has several aspirational moonshots, one being to pair each caregiver with a highly personalized AI assistant. The system also has strong aspirations to expand to broader smart/ambient hospital work.

"In healthcare, the cost of being wrong is very different, and we hold AI to that standard."

"We don't get to be wrong once and learn later."

"Our role is to make sure this is done the right way, not the fastest way."

Executive Perspective: Large Academic-Affiliated Health System(Continued)

2

Key Barriers to AI Transformation

The primary barriers to broader AI adoption are trust, governance, and cultural expectations, rather than technical capability. AI systems are held to a higher performance threshold than humans, particularly in clinical and patient-facing contexts. This standard reflects how slow the pace at which new capabilities can be operationalized.

Governance requirements are extensive and intentionally conservative, with multiple layers of review and approval required before AI tools can be deployed. These processes are designed to ensure accountability and transparency, but they also introduce friction when scaling AI across departments. In addition, clinician trust must be earned through demonstrated reliability and clear communication around limitations.

Operationally, workflow integration remains a challenge. Even where AI tools are available, integrating them into complex clinical environments without increasing burden requires careful redesign and coordination across teams.

3

Strategic Aspirations and Moonshots

The organization's long-term direction emphasizes measured expansion of AI support in operationally complex areas, rather than transformational autonomy. Future investment is focused on improving predictive accuracy, reliability, and integration of existing decision-support tools, particularly in areas such as throughput management, scheduling, and financial operations.

Ambient or agentic concepts are discussed cautiously and remain exploratory. There is no articulated goal to automate clinical decision-making or delegate authority to AI systems. Instead, leadership prioritizes using AI to reduce noise, surface relevant signals earlier, and support more consistent execution of existing care and operational models.

The institution views its role as setting a high bar for responsible AI use in healthcare, even if that means adopting new capabilities more slowly than peers.

Key Takeaways

1. **AI adoption is governed by trust and accountability** more than technological readiness.
2. **Human oversight is a deliberate and durable design choice**, especially in clinical contexts.
3. **AI progress is expected to be incremental,** with emphasis on validation, governance, and safety over autonomy

Executive Perspective: Academic Specialty Health System (Oncology-Focused)

System and Executive Description: This case study reflects a large, specialty cancer center. Digital, data, and AI capabilities are tightly integrated into clinical operations, research, and access management, with a strong focus on supporting multidisciplinary care teams. The executive interviewed holds enterprise responsibility for digital and clinical technology strategy and works closely with clinical leadership to align AI initiatives with care delivery, workforce sustainability, and patient experience in oncology settings.

$7b

In total operating revenue

Nationwide Presence

1

Current State of AI Strategy & Maturity

AI is positioned as a high-priority but deliberately governed capability, shaped by the system's hyper-specialized oncology environment and long-standing investment in curated data assets. The organization applies the same rigor and review processes used for other enterprise technologies, with additional scrutiny related to data access, storage, model training, and vendor exposure. AI adoption is framed as a strategic lever to support growth while maintaining capital discipline, rather than as a standalone innovation initiative.

Today, AI is delivering targeted impact across clinical operations, administrative domains, research, and technology operations. Ambient documentation remains a priority, given the high-acuity and high-complexity oncology setting. Beyond documentation, AI supports access and throughput by analyzing referral patterns, scheduling complexity, and care sequencing across the continuum. These applications surface bottlenecks and variation in access to specialty services, enabling earlier intervention. AI is also used to organize and summarize complex clinical information to support visit preparation and multidisciplinary care, without automating clinical decisions, and has been influential in research and clinical trial acceleration.

AI investments play a critical role in non-clinical domains. In business and administrative domains, AI is being used to examine workflows related to financial tracking, human-capital deployment, and operational efficiency as part of a broader effort to modernize the operating model. Additionally, AI is being applied internally within technology operations, including ServiceNow-related workflows, to augment ticket creation, knowledge management, and capacity within IT teams.

The organization is actively investing in applications that reduce friction across the care journey, particularly where coordination, information synthesis, and timing are critical. While interest in agentic AI is high, deployment is domain-dependent. Assistive and orchestration functions are prioritized in administrative and operational workflows, while autonomous or semi-autonomous clinical decision-making is approached conservatively given the complexity and rarity of cancer cases. Human oversight and existing medical practice governance structures remain essential.

"Our responsibility is to get this right before we get it everywhere."

"Our challenge isn't generating data; it's helping clinicians make sense of it across the care journey."

"If we can't explain exactly how it works and how it fails, we won't put it into practice."

Case Study: Academic Specialty Health System (Oncology-Focused) (Continued)

2

Key Barriers to AI Transformation

The most significant barriers to AI transformation are risk tolerance, clinical trust, and the consequences of error in oncology care. Given the stakes involved in cancer diagnosis and treatment, the organization maintains a very low tolerance for incorrect or misleading outputs. AI systems are expected to meet a high standard for accuracy, transparency, and reliability before being embedded into clinical workflows.

Governance and oversight requirements are correspondingly rigorous. AI tools are evaluated not only for performance, but for how they interact with clinical judgment, patient communication, and multidisciplinary care processes. Even when AI shows promise, deployment can be delayed if workflows cannot be redesigned to safely absorb the technology without increasing cognitive load or introducing ambiguity.

Workforce considerations also play a central role. Clinicians must trust that AI tools support, rather than disrupt, their relationship with patients. Any perception that AI interferes with clinical authority or patient trust is treated as a material risk.

3

Strategic Aspirations and Moonshots

The organization's long-term vision for AI is closely tied to a reimagining of the care environment itself, rather than incremental automation of existing tasks. Leadership describes a future in which technology becomes largely invisible, enabling clinicians to focus fully on patients while AI manages background complexity. This includes expanded use of ambient technologies, real-time information synthesis, and AI-enabled coordination across care teams. This posture towards ambient is already informing current/future facility planning efforts.

Agentic AI is viewed as a potential future capability for managing coordination and information flow across the cancer care journey, particularly in areas such as care navigation, follow-up management, and cross-setting coordination. However, there is no stated ambition to delegate diagnostic or treatment decisions to autonomous systems. Human judgment and accountability remain central to the organization's vision.

Rather than pursuing rapid deployment, the organization aims to set standards for responsible AI use in high-acuity specialty care, even if that means slower adoption. The emphasis is on long-term transformation of the operating model, not short-term efficiency gains.

Key Takeaways

1. **AI is a lever for redesigning care delivery** in highly complex, high-acuity environments.
2. **Specialties with low tolerance for error require a conservative approach to autonomy,** even with high ambition for long-term transformations.
3. In high-stakes clinical environments, **ambient, assistive, and coordination-focused AI act as precursors** to future agentic capabilities.

Case Study: Regional Health System 1

System and Executive Description: This case study reflects a regional integrated health system serving a largely rural and mixed urban population. Digital and AI capabilities are led through centralized IT and analytics functions, with close coordination across clinical, financial, and administrative teams. The executive interviewed holds responsibility for enterprise digital and analytics initiatives and focuses on deploying technology that fits naturally into existing workflows across hospitals, ambulatory sites, and system operations.

$5b

In total operating revenue

Present in primarily **1** state

1

Current State of AI Strategy & Maturity

AI was described as a core strategic priority with clear concentration in a small number of domains where the organization has already realized tangible value. The executive framed ambient documentation as the foundational inflection point that demonstrated AI's real-world utility to clinicians and catalyzed broader demand across the system. AI efforts were not positioned as speculative; instead, they were described as pragmatic, use-case driven, and increasingly tied to measurable operational and financial outcomes, particularly in revenue cycle and productivity.

The areas seeing the most immediate impact are clinical documentation, revenue cycle operations, and administrative workflows. Ambient documentation is used to reduce clinician documentation time and improve visit flow, with success measured by adoption and time savings rather than downstream financial metrics. A pilot for nursing ambient documentation is planned, with intent to automate flowsheets. In revenue cycle operations, AI is automating denials appeals and prior authorization workflows, producing the only tangible and meaningful *financial* ROI.

Beyond current deployments, the organization is actively investing in embedded, workflow-level AI tools that address capacity management and throughput optimization, particularly those integrated directly into core platforms such as the EHR, ERP, and productivity environments. Leadership views platform-native AI as a way to minimize change management burden and accelerate adoption. As part of this approach, the system has enabled broad access to generative AI tools for administrative and knowledge work/retrieval tasks, with guardrails in place for appropriate use.

Agentic AI is approached cautiously but practically. Rather than pursuing fully autonomous workflows, the organization is exploring task-level automation where AI can draft, route, or prepare work for human approval. These efforts are focused on administrative and clinically adjacent domains, where risk is lower and outcomes are easier to validate.

"We're less interested in moonshots than in making the basics work better every day."

"The risk isn't the technology—it's rolling something out that doesn't quite fit."

"If it doesn't fit naturally into how people already work, it won't scale."

Case Study: Regional Health System 1 (Continued)

2

Key Barriers to AI Transformation

The primary barriers to broader AI adoption are workflow integration and governance clarity, rather than technical readiness. Even when AI tools demonstrate value, deployment can stall if they introduce additional steps or cognitive load for clinicians and staff. As a result, leadership prioritizes solutions that are embedded directly into existing systems and workflows.

Governance is evolving alongside adoption. While guardrails exist for data use and patient-facing applications, the rapid pace of AI tool availability, particularly generative AI, creates pressure to clarify acceptable use and oversight responsibilities. Balancing speed with appropriate controls remains an ongoing challenge.

Workforce readiness also plays a role. While interest in AI tools is high, there is variability in comfort and skill levels across roles. Ensuring that staff understand what AI can and cannot do is viewed as essential to avoiding misuse and overreliance.

3

Strategic Aspirations and Moonshots

Longer-term aspirations focus on deeper integration of AI into core operational platforms rather than pursuing standalone or experimental solutions. Leadership envisions AI increasingly handling preparation, drafting, and prioritization tasks across administrative and clinically adjacent workflows, allowing staff to focus on judgment, and communication. That stated, the executive discussed potential voice-based and video-based agents for both internal and patient-facing interactions. Internally, these could support service desk requests such as password resets. Externally, they could assist with routine patient interactions or front-desk functions, particularly in roles with high turnover such as patient service representatives. These ideas were framed as plausible but not yet ready for broad deployment, with patient comfort and acceptance identified as the primary gating factor rather than technology maturity.

Agentic concepts today are considered primarily in the context of bounded automation, such as AI managing task queues, coordinating handoffs, or preparing documentation for review. Fully autonomous clinical decision-making is not part of the organization's stated vision. Instead, the emphasis is on incremental expansion of AI support as reliability and governance mature.

The system's long-term posture reflects a belief that sustained AI value will come from consistent execution and integration, rather than from high-risk, transformational bets.

Key Takeaways

1. **Speed to value and workflow fit drive AI adoption** rather than long-term transformation narratives.
2. **Systems favor platform-integrated AI tools** as a way to reduce change management burden and accelerate adoption.
3. Agentic AI adoption focuses on task prioritization and preparation **rather than full autonomy.**

Case Study: Regional Health System 2

System and Executive Description: This case study reflects a regional integrated health system. Digital and AI initiatives are led through a centralized digital and information services organization, with strong linkage to operational and financial leadership. The executive interviewed holds responsibility for enterprise digital strategy and execution and works across clinical, operational, and IT teams to support near-term performance improvement and system execution priorities.

$3b

In total operating revenue

Present in **1** state with spillover into surrounding state counties

1 Current State of AI Strategy & Maturity

AI is positioned as a strategic but highly pragmatic capability within a regional integrated health system facing near-term operational and financial pressure. The system's AI strategy is explicitly oriented toward immediate, measurable value, with leadership intentionally deprioritizing long-range or speculative transformation narratives. AI is framed as one of several levers, alongside access and experience, to stabilize performance and support execution in the current operating environment.

The areas generating the most impact today are workforce management, revenue cycle operations, and clinical documentation. Predictive models are used to support nursing staffing and skill-based redeployment across units and sites, helping operational leaders respond more dynamically to demand variability. In revenue cycle, AI is applied to denials management and documentation completeness, surfacing missing or inconsistent information earlier in the process to protect reimbursement. In home health, AI is used to ensure required documentation elements (e.g., OASIS assessments) are consistently completed, reducing downstream payment risk.

These applications are valued for their operational reliability and financial protection, not for automation of judgment. AI outputs inform decisions and prompt action, but execution remains human-led. Across use cases, leadership emphasizes near-term performance improvement over experimentation. When discussing readiness for agentic or autonomous AI, the executive was explicit that the organization is not yet fully ready, suggesting partial technical capability paired with unresolved organizational, cultural, or governance gaps. Movement beyond assistive AI was discussed cautiously, without clear thresholds or timelines for advancement.

"The risk isn't building the model—it's putting something into practice before people are ready for it.

"We're focused on what actually moves the needle right now, not what might matter five years from now.

"We're optimistic about where this could go—but we're not rushing ahead of ourselves."

Case Study: Regional Health System 2 (Continued)

2

Key Barriers to AI Transformation

The primary barriers to broader AI adoption are organizational readiness, governance clarity, and change management, rather than access to technology. While the system has invested in data infrastructure and external platforms, leadership acknowledges a gap between technical capability and the organization's ability to absorb more advanced AI into everyday workflows.

Trust and role clarity represent additional constraints, particularly in clinical contexts. Leadership draws a firm boundary around clinical authority, emphasizing that AI should inform but not direct care decisions. Ensuring that clinicians understand where AI fits—and where it does not—is viewed as essential to adoption. Change management is treated as a first-order requirement, with explicit attention to explaining the "why," involving staff in workflow redesign, and addressing concerns about role displacement.

Unlike some peers, clinician demand for AI is not described as outpacing deployment. Instead, leadership is intentionally pacing adoption to align with governance capacity and workforce readiness.

3

Strategic Aspirations and Moonshots

Longer-term aspirations reflect measured interest in agentic and more autonomous capabilities, but without a defined timeline or roadmap. Leadership expresses optimism about agentic AI's potential to manage complexity and reduce cognitive burden as data volume continues to grow. However, these discussions remain largely conceptual and are framed as future possibilities rather than near-term priorities.

Speculative ideas include more advanced automation in administrative and documentation-heavy workflows, as well as exploratory concepts such as digital or avatar-based clinical representations. These ideas are discussed cautiously and with recognition that cultural acceptance, trust, and governance would need to mature significantly before such capabilities could be realized.

Overall, the system's strategic posture favors execution discipline over transformational risk, with innovation progressing incrementally as confidence, capability, and alignment improve.

Key Takeaways

1. **AI delivers the most value as a near-term operational and financial stabilizer** rather than a long-horizon transformation engine.
2. **AI effectively informs staffing, documentation, and revenue protection decisions,** with humans retaining execution and accountability.
3. Interest in agentic AI exists, but **organizational readiness and change management determine the pace of progress.**

Research Methodology and Support

In 2025, The Health Management Academy administered a quantitative survey to Leading Health Systems executives regarding their perspectives on AI transformation, and agentic AI specifically. This survey was supported by several qualitative interviews conducted with AI leaders at health systems to further contextualize findings. The 30 quantitative survey responses represent 30 total unique executives from 30 unique health systems.

The Health Management Academy Project Team

- **Coltin Ball,** Senior Analyst, Industry Insights
- **Emma Feeney**, Associate Director, Industry Insights
- **Wes Adams,** Senior Director, Industry Insights

About The Health Management Academy

The Health Management Academy powers our community to drive health forward. Our community is made up of healthcare's most influential changemakers including executives from the top 150 U.S. health systems and the most innovative industry partners. We power our members by building our community and fostering connections through executive peer learning. We support professional growth through talent and development. We accelerate understanding by delivering timely and actionable data and insights on key challenges. And we catalyze transformation by building alliances in areas where the power of the collective is greater than the power of one.

More information is available at www.hmacademy.com.