


# The intelligent health system: How AI is redefining healthcare operations



# Table of contents

01. Introduction: The operational imperative	3
02. AI in healthcare: From automation to intelligent operations	4
03. Why healthcare has been cautious, and why that's changing	5
04. Autonomous agents: A digital extension of the healthcare workforce	6
05. Smarter financial operations and revenue cycle management	7
06. Intelligent supply chain and procurement	8
07. Demand-driven workforce management	9
08. Access, throughput, and patient flow	10
09. Unified data: The foundation of operational intelligence	11
10. Cloud-ready infrastructure: Enabling scale and speed	12
11. AI as a force multiplier, with humans always in control	13
12. Conclusion: Building the intelligent health system	14

# Introduction: The operational imperative

Healthcare organizations are caught between two forces that show no sign of easing. Patient demand continues to rise while the available workforce is shrinking. Physician shortages could reach 86,000 by 2036, nursing turnover remains elevated, and clinicians spend nearly two hours on paperwork and EHR tasks for every hour of direct care<sup>1, 2</sup>. At the same time, operating costs have risen sharply, and many health systems are operating at or near breakeven.

To modernize operations and reduce the burden on staff, healthcare organizations have invested heavily in digital transformation over the past decade. EHRs, telehealth platforms, and remote monitoring tools have expanded access and modernized elements of care delivery. Yet these investments largely digitized existing processes rather than changing the operating model. The technology stack grew, but the underlying workflows of scheduling, procurement, staffing, revenue cycle, and supply chain remained labor-intensive and reactive.

Artificial intelligence (AI) represents a different kind of opportunity. Not as another system to learn or another platform to manage, but as an intelligence layer that works within the systems organizations already use. When applied to operational and administrative workflows, AI can process information, coordinate tasks, and support decisions at a speed and scale that human teams alone cannot sustain. The value is not in replacing clinical judgment. It is in removing the operational friction that keeps clinicians and staff from focusing on what they do best: caring for patients.

While there is significant excitement around AI in healthcare, adoption has remained cautious and incremental, concentrated in pilots rather than embedded in operations. The organizations that move beyond experimentation to operational deployment will be best positioned to deliver quality, control costs, and retain talent in the years ahead.

---

<sup>1</sup> Association of American Medical Colleges. (2024). *The complexities of physician supply and demand: Projections from 2021 to 2036*.

<sup>2</sup> Sinsky, C. A., Colligan, L., Li, L., Prgomet, M., Reynolds, S., Goeders, L., Westbrook, J. I., Tutty, M., & Blike, G. (2016). Allocation of physician time in ambulatory practice: A time and motion study in 4 specialties. *Annals of Internal Medicine*.

# AI in healthcare: From automation to intelligent operations

Healthcare's relationship with automation is not new. Rules-based systems have long handled tasks like eligibility verification, claims edits, and appointment reminders. These tools delivered value but operated in narrow lanes, each one addressing a single task and requiring people to stitch activities together across systems.

Generative AI introduced a more capable layer. In Q4 2025, McKinsey found that the majority of healthcare organizations had either implemented generative AI use cases or begun developing proofs of concept, signaling rapid maturation<sup>3</sup>. These systems can interpret unstructured data, summarize complex records, generate content, and surface contextual recommendations.

Now a third phase is emerging: agentic AI. Where generative AI creates and summarizes, agents act. These agents can operate with varying degrees of autonomy, from human-in-the-loop systems that require a person to initiate or approve each step, to fully autonomous agents that complete end-to-end workflows within defined guardrails. They can initiate and coordinate multi-step workflows that span departments and platforms, such as triggering procurement actions when supply levels shift. Rather than waiting for a human to interpret a recommendation and take action, these systems execute automatically, escalating to people only when judgment or approval is required.

Critically, this is not about introducing clinical AI that makes diagnostic or treatment decisions. The most immediate and impactful opportunity lies in the operational and administrative domains: finance, supply chain, workforce management, and revenue cycle. This matters now more than ever, as health systems operate on tighter margins and face persistent workforce shortages that demand new ways to do more with less.



<sup>3</sup>Lamb, J., Dahab, A., & Jones, J. (2026, April 16). *Generative AI in healthcare: Adoption matures as agentic AI emerges*. McKinsey & Company. <https://www.mckinsey.com/industries/healthcare/our-insights/generative-ai-in-healthcare-current-trends-and-future-outlook>

# Why healthcare has been cautious, and why that's changing

Health system leaders have good reasons to approach AI carefully. The consequences of error in this industry are measured in ways that other sectors rarely face, and the cost of getting it wrong can outweigh the gains of moving fast. But the question facing health systems today is no longer whether to adopt AI, it is how. The answer is a practical one, grounded in three principles gaining traction across the industry.

**1**

**AI should be embedded within the systems organizations already use**, including enterprise resource planning (ERP), workforce management platform, and financial systems, rather than layered on as a standalone tool. This approach keeps data inside trusted environments and ensures AI can act on the same information that drives day-to-day operations.

**2**

**AI should start with operational and administrative use cases where the risk profile is lower** and the efficiency gains are immediate. Starting with repeatable, structured work in finance, supply chain, workforce management, and revenue cycle—far removed from the life-or-death decisions of clinical care—lets organizations build confidence before extending AI into more sensitive domains.

**3**

**Every AI-driven process should maintain human oversight**, with clear mechanisms for people to review, approve, or override actions at critical decision points. This oversight should be reinforced by a formal AI governance plan that defines where AI is permitted, how it is evaluated, and who is accountable. Established frameworks like the National Institute of Standards and Technology (NIST) AI Risk Management Framework offer health systems a practical, widely adopted starting point.

Organizations that adopt this approach are finding that AI does not add complexity. It reduces it. When an AI agent processes non-purchase order (PO) invoices, validates them, and routes exceptions for review, the procurement team gains capacity rather than another system to learn. When a scheduling system fills an open shift before a manager makes a single phone call, that manager gets time back. Well-designed AI does not create more work. It absorbs the operational burden that has been steadily growing for years.

# Autonomous agents: A digital extension of the healthcare workforce

Workforce shortages are not a future risk. They are a present reality affecting virtually every health system in the country. At the same time, a significant share of staff time is consumed by administrative and operational tasks that do not directly contribute to patient care. For example, nurses spend hours on documentation and coordination, while finance teams manually reconcile invoices.

Autonomous agents represent a shift from AI that assists to AI that executes. Unlike traditional automation, which follows fixed rules, these systems adapt to changing conditions, manage dependencies across multiple steps, and coordinate actions across systems. This makes them especially effective for workflows that span departments, involve multiple systems, and require coordination under time pressure.

Consider a practical example. When a patient is identified as ready for discharge, the process involves clinical documentation, pharmacy reconciliation, transport coordination, bed management notification, follow-up scheduling, and insurance communication. Today, much of this is coordinated manually through phone calls, messages, and task lists, creating delays that ripple across the organization. An agentic system can orchestrate these steps in parallel, triggering each action as prerequisites are met and escalating to staff only when a decision requires human judgment.

Deloitte has noted that Mayo Clinic is deploying AI agent approaches to streamline administrative workflows between providers and payers, including eligibility verification, prior authorization, utilization management, and prescription support<sup>4</sup>. By reducing manual touchpoints, these systems reclaim time for clinical teams and accelerate payer coordination.

The goal is to extend capacity, not eliminate staff. The future of work will be hybrid: human teams supported by AI systems that handle routine operational tasks so that people can focus on work requiring empathy, clinical expertise, and professional judgment. And the timing matters: Deloitte research finds that 59% of early adopters expect cost savings above 20% in the next two to three years, versus just 13% of watchers<sup>4</sup>.

<sup>4</sup>Bruno, J., Fera, B., Yung, J., McCallen, M., Bhatt, J., & Shukla, M. (2026, February 11). *Many health care leaders are leaning into agentic AI as adoption hurdles ease*. Deloitte Center for Health Solutions. <https://www.deloitte.com/us/en/insights/industry/health-care/agentic-ai-health-care-operating-model-change.html>

# Smarter financial operations and revenue cycle management

Revenue cycle management is one of the most labor-intensive and error-prone domains in healthcare operations. From charge capture and coding to claims submission, denial management, and payment reconciliation, the process involves dozens of handoffs, complex payer rules, and large volumes of transactions.

AI is beginning to transform this domain by automating the repetitive, high-volume steps that consume staff time. Systems can predict which claims are likely to be denied before submission, flag coding inconsistencies, and automate the initial review of remittance data. Cash flow forecasting models can analyze historical patterns, payer behavior, and seasonal trends to give finance leaders more accurate and timely visibility into expected revenue.

On the payables side, healthcare procurement creates unique challenges. Unlike manufacturing or distribution, where invoices typically map to purchase orders, healthcare invoices are frequently non-PO based, arriving via email, fax, or supplier portals. They often include unique device identifiers, medical procedure codes, and patient-adjacent references that require specialized validation. Matching is non-standard, often service-based or patient-based, and compliance scrutiny is high because of the connection to reimbursement and regulatory requirements.



AI agents can parse these documents, extract relevant fields, validate them against contract terms and compliance rules, and route exceptions for human review, dramatically reducing the manual burden on accounts payable teams while improving accuracy and audit readiness. When combined with general ledger automation and cash operations intelligence, these capabilities give finance leaders a more complete, real-time view of organizational financial health.



The opportunity also extends to financial planning and analysis. Traditional healthcare reporting often relies on manually compiled reports that take days to produce. AI-powered analytics can allow finance leaders to ask natural-language questions and receive answers in seconds, grounded in real-time data from across the organization. This shift from static reporting to dynamic intelligence enables faster, more confident decision-making.

# Intelligent supply chain and procurement

Healthcare supply chains are among the most complex in any industry. A typical health system manages thousands of vendors, millions of purchase orders and invoices, and inventory ranging from surgical implants to everyday clinical supplies, all while navigating group purchasing agreements, regulatory requirements, and constant disruption risk.

Yet procurement teams still spend disproportionate time on administrative tasks like processing requisitions and matching invoices, rather than on strategic sourcing and cost management. AI introduces the ability to shift from reactive to proactive supply chain management.

Demand forecasting models can analyze historical consumption patterns alongside external signals like seasonal trends, scheduled surgical volumes, and census projections to predict what supplies will be needed, where, and when. Reorder points can be dynamically optimized rather than set statically, reducing both stockouts and excess inventory. Vendor performance data can be systematically analyzed to identify reliability issues, late deliveries, or inconsistent fill rates, giving procurement teams the intelligence they need to make better sourcing decisions.

For invoice processing specifically, AI can address one of the most persistent pain points in healthcare procurement. Healthcare invoices frequently arrive in non-standard formats, include specialized codes, and require compliance validation that differs from other industries. As a result, traditional automation has struggled to handle them. Newer AI systems can extract structured data from unstructured documents, validate it against contract terms and regulatory requirements, match it against receiving records, and flag anomalies, all while maintaining the audit trail that healthcare compliance demands.

The cumulative impact is significant. When supply chain teams spend less time on manual processing and more time on strategic decisions, organizations reduce costs, minimize waste, and improve their ability to respond to disruption. When supply chain challenges can directly affect patient care, as the pandemic starkly demonstrated, these capabilities are not just operational improvements. They are resilience measures.

# Demand-driven workforce management

Workforce management is arguably the single most consequential operational challenge in healthcare today. Labor represents the largest line item in most health system budgets, typically 50% or more of total operating expenses, and the combination of persistent shortages, high turnover, and rising agency staffing costs has made it an urgent strategic priority<sup>5</sup>.

Traditional scheduling approaches compound the problem. When a nurse calls out before a shift, a charge nurse or manager still spends an hour or more manually working through availability, labor rules, and credentialing, often leaving gaps.

AI-enabled workforce management offers a fundamentally different model. By analyzing real-time census data, acuity levels, historical demand patterns, and staffing constraints, systems can predict what each unit will need before the schedule is built. When gaps emerge, intelligent scheduling tools can identify qualified, available staff based on credentials, labor rules, overtime thresholds, and employee preferences. Open shifts can be broadcast to eligible employees, who can claim them through a simple mobile interface, filling coverage gaps in minutes rather than hours.

Labor forecasting extends this further. By modeling expected patient volumes, seasonal patterns, and planned procedures against current staffing levels, organizations can anticipate shortfalls weeks in advance and adjust hiring, float pool deployment, or agency utilization proactively.



When combined with attendance tracking that can surface anomalies and patterns, such as units with chronic callout issues or time-and-attendance discrepancies, leaders gain the visibility they need to address workforce challenges systemically rather than reactively.

The impact goes beyond efficiency. When staffing more closely matches patient needs, burnout decreases, care continuity improves, and patients receive more timely attention. When managers spend less time filling shifts and more time supporting their teams, unit culture and retention improve.

<sup>5</sup>American Hospital Association. (2026, March). *Costs of caring: Challenges facing America's hospitals as they care for patients in 2026*. <https://www.aha.org/costsofcaring>

# Access, throughput, and patient flow

Access to care has become a defining test of health system performance. Long wait times for appointments, emergency department boarding, surgical scheduling backlogs, and delayed discharges all contribute to a patient experience that often falls short of expectations, even when clinical quality is high.

These bottlenecks are rarely caused by a lack of resources in absolute terms. More often, they result from inefficiencies in how resources are coordinated. Beds remain occupied hours after a patient is ready for discharge because downstream steps have not been initiated. Operating rooms sit idle between cases because turnover processes are not optimized. Outpatient appointments are scheduled based on rigid templates that do not reflect actual demand patterns.

AI addresses these challenges by improving coordination and prediction. Demand forecasting can anticipate patient volumes by department and time period,

allowing organizations to align staffing and capacity more effectively. Process mining can analyze actual patient flow, from emergency department (ED) arrival to discharge, identifying where delays occur. Predictive models can flag patients at risk of extended stays, enabling proactive intervention.

When these capabilities are integrated with scheduling, bed management, and staffing systems, resources are deployed where needed, bottlenecks are addressed before they create backups, and throughput improves without requiring additional capacity.



# Unified data: The foundation of operational intelligence

AI models are only as good as their data<sup>6</sup>. In healthcare, data is notoriously fragmented: clinical information lives in the EHR, financial data in the ERP, supply chain data in procurement systems, workforce data in scheduling and HR platforms, and operational metrics in a patchwork of departmental tools and spreadsheets. Inconsistent or incomplete data leads to unreliable outputs.

This fragmentation does not just limit AI, it limits the organization. When staffing decisions cannot account for real-time patient demand because census data lives in a different system, the result is misalignment. When supply chain teams cannot see utilization patterns because data is locked in clinical systems, the result is waste. When finance leaders cannot connect operational performance to financial outcomes in real time, decisions are delayed and incomplete.

Unified data environments address these challenges by connecting information across systems and creating shared context. When clinical, operational, and financial data flow through a common platform with consistent definitions and real-time access, AI can generate insights that reflect the full picture. Forecasting becomes more accurate because it draws on more complete inputs and agents can coordinate across domains because they share access to the same underlying data.

Building this foundation is not trivial, and it does not happen overnight. But it is the single most important enabler of AI at scale. Organizations that invest in connected, well-governed data platforms will deploy AI capabilities faster and achieve better results as technology continues to evolve.



<sup>6</sup> Jonker, A., & Aquino, J. (n.d.). *Why AI data quality is the key to AI success*. IBM. <https://www.ibm.com/think/topics/ai-data-quality>

# Cloud-ready infrastructure: Enabling scale and speed

Legacy on-premises systems were not designed for the speed, scale, and integration that AI requires. They often operate in silos, with limited ability to share data across applications or support real-time processing.

Cloud-based platforms address these limitations by enabling real-time data access, seamless integration across systems, and the ability to scale resources as needed. For healthcare organizations, cloud adoption is becoming a strategic priority not only for IT modernization but also as the foundation for AI-driven operations. When AI agents, workforce management tools, financial systems, and supply chain platforms all operate on a connected cloud platform, the barriers to coordination are reduced.

Information flows where it is needed, actions can be triggered across systems, and the organization can move faster without sacrificing control.

Modern cloud platforms are also designed with the security, compliance, and governance requirements that healthcare demands. Role-based access controls, encryption, audit logging, and regulatory compliance frameworks are built into the infrastructure, ensuring that the shift to the cloud does not compromise the protections healthcare organizations require.



# AI as a force multiplier, with humans always in control

The most important principle in healthcare AI adoption is also the simplest: AI should make people more effective, not replace them. In an industry built on human expertise, empathy, and trust, technology succeeds when it amplifies these qualities rather than attempting to substitute for them.

This principle means designing AI systems with human oversight at every critical decision point. When an AI agent processes a batch of invoices, a human reviews the exceptions and approves final postings. When a scheduling system fills an open shift, a manager can review and override the assignment. When a forecasting model predicts a surge in demand, leaders decide how to respond. The AI handles the volume, the speed, and the coordination. The human provides the judgment, the context, and the accountability.

Transparency is equally essential. Healthcare professionals need to understand how an AI system arrived at a recommendation or action. Systems that provide visibility into their reasoning process—showing which data informed a decision and what alternatives were considered—earn trust far more effectively than tools that simply deliver outputs.



Governance frameworks must be clear and enforceable: defining which processes AI can touch and which remain human-led, establishing role-based access so only authorized users interact with specific AI capabilities, maintaining audit trails for every AI-assisted action, and ensuring compliance with the Health Insurance Portability and Accountability Act (HIPAA) and other regulatory requirements at every layer. When these safeguards are in place, AI becomes what healthcare organizations actually need: a force multiplier that absorbs administrative and operational burden so that clinicians, nurses, finance professionals, and operational leaders can direct their expertise where it matters most.



# Conclusion: Building the intelligent health system

The pressures facing healthcare, including workforce shortages, rising costs, growing demand, and regulatory complexity, are not cyclical. They are structural. The traditional playbook of incremental efficiency gains and periodic technology upgrades is no longer sufficient to address them.

AI offers a fundamentally different approach: the ability to embed intelligence into the operational fabric of the organization, enabling systems that anticipate needs, coordinate workflows, and execute routine tasks at a speed and scale that human teams cannot match alone—not as a replacement for people, but as the operational infrastructure that allows people to focus on the highest-value work.

To capture this value, healthcare leaders should focus on four priorities:



## Establish a clear AI usage plan

Define how AI will be applied across the business and classify use cases by the level of autonomy and risk they carry. A useful framework organizes AI into four categories, each requiring different controls and oversight:

- Informational AI (FAQs, summaries)
- Analytical AI (forecasting, detection)
- Advisory AI (recommendations)
- Operational AI (execution)





### Invest in connected data and modern platforms

AI cannot operate reliably without a foundation of unified, high-quality, real-time data. Organizations that build this foundation now will deploy AI capabilities faster and achieve better results as the technology matures.



### Prioritize high-impact operational use cases

The greatest near-term value lies in supply chain optimization, workforce management, revenue cycle automation, and patient flow coordination, where workflows are repeatable and efficiency gains translate directly to financial performance.



### Establish clear governance with human oversight at the center

AI earns trust through transparency, accountability, and control. Organizations that define where AI acts, where humans decide, and how every action is tracked will scale faster and more confidently than those that treat governance as an afterthought.

The future of healthcare operations will not be defined by organizations that adopt the most technology. It will be defined by organizations that embed intelligence into how they operate in a thoughtful, strategic way, always in service of the people who deliver care and the patients who receive it.



## About Infor

Infor is a global leader in business cloud software products for companies in industry-specific markets. Infor builds complete industry suites in the cloud and efficiently deploys technology that puts the user experience first, leverages data science, and integrates easily into existing systems. Over 67,000 organizations worldwide rely on Infor to help overcome market disruptions and achieve business-wide digital transformation.

[infor.com](https://www.infor.com)

**Modernize  
healthcare operations  
with Infor Velocity Suite**

[LEARN MORE](#)

Copyright © 2026 Infor. All rights reserved. The word and design marks set forth herein are trademarks and/or registered trademarks of Infor and/or related affiliates and subsidiaries. All other trademarks listed herein are the property of their respective owners.  
WP-EN-0526-86ah2uguv-1

**infor**