

Healthcare AI Readiness Framework for Health Systems

How health systems can prepare for AI that works in everyday operations



Whitepaper



Contents

1. Executive summary	3
2. Understanding AI readiness in healthcare	4
3. Why existing AI strategies are not delivering results.....	5
4. Pillar one: Data foundation readiness.....	6
5. Pillar two: AI governance readiness.....	7
6. Pillar three: Operational readiness	8
7. Pillar four: Agentic AI readiness.....	10
8. Readiness across healthcare segments.....	11
9. Operational readiness shapes AI success.....	12



Executive Summary

Why do healthcare AI initiatives struggle to become part of everyday operations?

A pattern we keep seeing in our conversations with healthcare professionals is that they are already pumped with too many ideas.

They clearly understand the kind of opportunities AI can bring to their operations. Even then, there are certain challenges that still pull them away from being completely AI-ready.

AI readiness in healthcare is not just about technology infrastructure or model performance.

It comes down to whether an organization has the data engineering practices, AI governance strategy, workflows, and operational foundations needed to support healthcare initiatives.

Readiness doesn't look the same everywhere.

- A provider network has different concerns than a payer.
- The same goes for a diagnostics organization, as it faces different challenges than a digital health platform.

Each segment is working toward the same goal, but the path to getting there can look very different. We'll further explore where organizations typically run into challenges, how those challenges differ across healthcare segments, and what it takes to move AI initiatives from promising pilots into everyday operations.




Understanding AI Readiness in Healthcare

A 2025 analysis^[1] found that 95% of enterprise AI pilots failed to generate measurable ROI. Quite some time back, when people spoke about AI readiness, the discussions always centered around whether there was enough data, cloud infrastructure, and integration bandwidth.

Right now, these are not priorities. The way AI readiness is looked at today has changed. It needs to be a bigger part of clinical documentation, patient engagement, claims operations, diagnostics workflows, and care management.

Strategy and adoption are usually confused with readiness, while all three serve different purposes. A strategy answers where an organization wants to go.

Adoption focuses on implementing tools and use cases. Readiness determines whether those initiatives can operate reliably once they become part of everyday business processes.





Why is the old definition of readiness no longer enough? Because before implementing AI, healthcare leaders need to answer a few things, such as:

1. How much trust is there in the data being used?
2. Will existing workflows support AI without creating additional complexity?
3. Are governance processes equipped to scale as AI expands across departments?
4. What mechanisms are in place to monitor AI systems after they move beyond controlled pilots?

Why Existing AI Strategies Are Not Delivering Results

Once an AI initiative gets approved, the excitement is usually high in the beginning, to the extent that even pilot programs are launched faster. The challenges of implementing AI in healthcare show up right after this, when questions about ownership begin to surface. This is what we refer to as the AI execution gap.

The distance between an approved AI initiative and an operational AI system that consistently delivers value is significant. One of the biggest contributors to this gap is workflow disconnect.

AI business solutions are often introduced into existing processes. A recommendation may be accurate, but if it arrives at the wrong point in a clinician's workflow, it becomes another notification to ignore. A similar challenge can be seen in claims operations, where teams are expected to adopt AI-driven decisions without changes to the surrounding processes.

Ownership turns out to be the next issue. When the AI solution is set up and running, if there is no designated individual to own performance monitoring or decide when changes are needed, it will struggle to move forward.




Data fragmentation adds another layer of complexity. Healthtech companies generate enormous amounts of data, but it often sits across different systems and departments.

Governance is often treated as a later-stage discussion, where companies find themselves playing catch-up. Questions around compliance, validation, auditability, and risk management tend to grow as AI moves closer to operational use.

In some cases, governance discussions begin only after a solution is ready for deployment, creating delays that could have been addressed much earlier.

AI Momentum Declines at Every Stage



Every AI solution in the market is meant to reduce manual effort. That is not the same as creating business value for providers, payers, diagnostics organizations, and digital health companies.

Pillar One: Data Foundation Readiness

Conversations about AI implementation in healthcare eventually circle back to data. Our experts have spent a considerable amount of time finding out the main reason why it happens. We've seen that AI struggles to see the full picture. Part of the challenge is that more than 80% of healthcare information^[2] exists as unstructured data.



As a patient goes through different procedures and treatments, their data also moves through multiple systems across departments. Bits of information are scattered across EHRs, claims systems, diagnostic platforms, and connected devices.

Bringing all of that together in a way AI can understand is often harder than building the model itself. Take these into consideration to understand whether your data can support AI and deliver value.

Data Accessibility	Data Interoperability	Data Lineage	Data Reliability
Information should be available when it is needed, rather than being locked behind manual extraction processes or dependent on a small group of specialists.	Healthcare providers work across EHRs, imaging platforms, and laboratory systems.	Leaders need visibility into where information originated and how it was used.	AI models should not encounter one version of the data during testing and a different version after deployment.
It takes a lot of time to locate, prepare, and move data before AI can use it.	Payers rely on claims, member, and provider data. Remote monitoring programs continuously generate information outside traditional clinical environments.	If an AI-generated recommendation cannot be traced back to its source, trust becomes difficult to establish.	Standardized structures and dependable data flows help AI perform as expected everywhere it is used.
A lot of AI challenges start disappearing when teams can access the data they need without delays.	Readiness depends on whether these systems can exchange information in a way that preserves meaning and context.	Data lineage provides the transparency needed to support AI governance strategy, validation, and accountability.	Reliability becomes especially important when decisions affect patient care, claims outcomes, operational efficiency, or regulatory reporting.
	Connecting systems is important. Ensuring they speak the same language is what allows AI to generate reliable insights.		

Pillar Two: AI Governance Readiness

Healthtech companies aren't starting from scratch when it comes to governance. When AI models enter the picture, things become a little different. As these systems learn and evolve, many teams are still figuring out how to manage their impact on clinical and administrative decisions.

Companies find themselves managing AI through informal processes that weren't designed for it. This is where an **AI accountability model** becomes important.

It establishes ownership and decision-making responsibilities throughout the lifecycle of an AI system. Building that level of accountability requires a few AI data governance capabilities.

 Governance Structure	• We need to have a designated person to own the decisions and risks. Right from approving AI use cases and validating models to monitoring performance and deciding what happens when something goes wrong, ownership is a must. • The goal is to make sure accountability is clear from day one.
 Model Lifecycle Management	• Deploying a model is not the finish line, and environments where AI systems operate constantly change. How data enters the system, real user behavior changes based on time. • Model performance is not something you can observe well after the fact. A well-performing model can drift away from its original expectations without validation, rare performance reviews. • Treating model lifecycle management as a continuous process helps you maintain oversight and ensure models remain compliant with regulatory requirements.
 Explainability Requirements	• The more influence AI has on decisions, the more important explainability becomes. • Users need to understand why a recommendation was made, especially when the outcome affects patient care. It is difficult to gain trust when AI recommendations cannot be explained or supported by evidence.
 Risk Evaluation Processes	• AI implementation does not guarantee perfect outcomes under every condition. • You might get unexpected outputs when the conditions in production differ from those used during development. Thus, we should be clear guidelines for escalating risks and determining when human intervention is required.
 Regulatory Alignment	• Healthcare providers and payers do not need to predict every future regulation. They do need governance structures that can adapt as expectations change. • When new requirements emerge, organizations can show they have already started by addressing compliance gaps or purpose AI models with governance built in from the start. • When models help make critical decisions against patients, data transparency and evidence safety. Strong governance helps organizations respond to regulatory changes proactively rather than reactively.

Pillar Three: Operational Readiness

Deployment failures has nothing to do with the model. An AI model can generate recommendations all day long, but that does not automatically improve patient care.

For example, a claims reviewer may receive an AI-generated recommendation without enough context to make a decision. People, processes, and responsibilities ultimately determine whether those recommendations create value. Operational readiness focuses on building that connection.

- **Establish a Clear Path from Recommendation to Action**

Define what happens next before deployment begins. Every recommendation should trigger a known action and lead to an expected outcome.

If a clear path to action is unavailable, recommendations may be ignored, delayed, or applied inconsistently.



- **Clarify Where AI Supports Decisions**

AI does not need to participate in every decision. Some decisions benefit from AI support and others require human judgment and regulatory oversight.

Clear operational boundaries improve readiness by defining how AI implementation supports decisions and who holds final decision-making authority.

- **Assign Operational Responsibilities Early**

Every healthcare professional who interacts with an AI-enabled solution should understand their role in reviewing outputs, validating recommendations, escalating concerns, and taking action when needed.

From an AI governance perspective, accountability should remain with people, not the technology.

- **Measure Outcomes Instead of Activity**

System usage, model performance, and adoption metrics are the kinds of metrics that only show whether AI is being used and do not indicate how AI is creating value.

Operational readiness measures help you look at the outcomes that matter to the business. Tie success metrics to operational outcomes, and it becomes easier to understand the real-world impact of AI.

Expert's Note

It is common practice to focus on reviewing situations where AI produces an incorrect recommendation. Operationally mature organizations review both sides of the equation.

They investigate when AI is wrong, and they also examine situations where AI was correct and no action was taken. Looking at both scenarios often reveals insights that are just as valuable as improving the model itself.



Pillar Four: Agentic AI Readiness

The benefits of generative AI^[4] are becoming easier to see across healthcare through new approaches to supporting patient care. Applications such as conversational AI in healthcare have shown how AI can improve access to information and make interactions easier for care teams.

Agentic AI builds on that foundation by moving beyond information generation and toward autonomous execution.

To prepare for that shift, healthcare leaders should focus on three capabilities.

 Set Autonomy Boundaries	• One of the biggest misconceptions about agentic AI is that autonomy is an all-or-nothing decision. In practice, autonomy exists on a spectrum. • Some actions can be executed independently with minimal risk, whereas others require approval and human validation before execution. • The objective here is to ensure autonomy operates within clearly defined boundaries.
 Build Multi-Layer Escalation Frameworks	• Healthcare workflows do not follow predictable paths because patient conditions change and priorities shift. • Agentic systems need a structured way to respond when operating conditions move outside predefined thresholds. • Build escalation frameworks that allow agents to transfer responsibility based on confidence levels, exception severity, business rules, or clinical risk. • The framework should also account for situations involving sensitive patient information and Personally Identifiable Information (PII). • Agents may encounter information they are not authorized to access or act on. In such situations, responsibility should be transferred to the appropriate team or individual for review. • This creates a controlled operating model where autonomy expands without sacrificing oversight.
 Create Permission-Based Execution Models	• When you deploy multiple agents, permission management becomes important. • An agent responsible for appointment coordination should not have the same authority as an agent supporting utilization management or care interventions. • Readiness requires organizations to establish permission structures based on responsibilities and workflow requirements.



Readiness Across Healthcare Segments

Each healthcare segment reaches AI maturity through a different path. The challenge is preparing the environment in which AI will operate.

Segment	Primary Readiness Challenge	Readiness in Practice
 Providers	Integrating AI into clinical workflows without increasing cognitive burden on care teams.	AI supports clinical decisions without disrupting care delivery.
 Payers	Connecting AI recommendations to utilization management and claims decision processes.	Recommendations move through claims workflows consistently and efficiently.
 Diagnostics	Establishing trust in AI outputs through validation, explainability, and review processes.	AI findings can be validated and trusted before decisions are made.
 Telehealth	Maintaining context across fragmented patient interactions and distributed care environments.	AI maintains a complete view of the patient journey across channels.
 Remote Monitoring	Creating interoperability across devices, platforms, and escalation pathways.	Patient alerts move quickly from detection to intervention.
 Digital Therapeutics	Measuring clinical outcomes consistently enough to support AI-driven interventions.	Outcomes can be tracked and linked to interventions over time.
 Patient Experience	Balancing personalization with privacy, trust, and engagement expectations.	AI delivers personalized experiences without compromising trust.
 MedTech	Aligning innovation cycles with evolving regulatory and safety requirements.	AI innovation scales alongside safety and regulatory requirements.



Segment vs Readiness Pillars Matrix

● Low Readiness ● Moderate Readiness ● High Readiness

Healthcare Segments	Data	Governance	Operations	Agentic
 Health Systems & Hospitals	● Moderate Siloed data across systems	● Moderate Policies exist but inconsistent	● High Mature workflows and digitization	● Moderate Early pilots in select use cases
 Physician Practices	● Low Limited data standardization	● Low Minimal governance framework	● Moderate Basic automation in place	● Low Exploring agentic capabilities
 Payers	● High Strong data infrastructure	● Moderate Well-defined policies and controls	● High Advanced process automation	● Moderate Use cases in development
 Life Sciences	● High Rich, standardized datasets	● High Strong compliance and governance	● Moderate R&D and commercial alignment in progress	● Moderate Active experimentation across functions
 Ambulatory & Specialty Providers	● Low Fragmented data sources	● Low Limited governance maturity	● Moderate Operational improvements underway	● Low Early stage exploration



How to read: Readiness levels indicate how prepared each segment is across the four pillars. Use this matrix to identify strengths, gaps, and priority areas for your organization.

Operational Readiness Shapes AI Success

An industry research study by Deloitte^[3] found that 61% of healthcare leaders are already building agentic AI initiatives or have secured budgets, while 85% expect investment to increase over the next few years.

This shows how AI will continue to become more capable and take on greater operational responsibility. Readiness is becoming the differentiator because it determines whether those initiatives can positively impact clinical^[5] operations and patient engagement.

AI has enough people talking about what it can do. We focus on what it needs to keep working through data engineering, AI governance, and agentic AI capabilities.

Every healthcare organization has
a different path to AI readiness.

Let's identify yours.

Connect with our experts →



References

1. Estrada, S. (2025). MIT report: 95% of generative AI pilots at companies are failing. Fortune. Available at: [MIT report: 95% of generative AI pilots at companies are failing | Fortune](#) [Accessed 20 Aug. 2026].
2. Kalankesh, L.R. and Elham Monaghesh (2024). Utilization of EHRs for clinical trials: a systematic review. BMC Medical research methodology, doi:10.1186/s12874-024-02177-7.
3. Many health care leaders are leaning into agentic AI as adoption hurdles ease. (2026). Deloitte Insights. Available at: [Many health care leaders are leaning into agentic AI as adoption hurdles ease](#) [Accessed 20 Aug. 2026].
4. The Power of Predictive Analytics in Healthcare: Using Generative AI and Confluent. (2024). Confluent. Available at: [The Power of Predictive Analytics in Healthcare: Using Generative AI and Confluent](#) [Accessed 20 Aug. 2026].
5. Pollen, A. (2024). Measuring True Readiness for Clinical AI in Health Systems | Aidoc. Healthcare AI | Aidoc Always-on AI. Available at: [Measuring True Readiness for Clinical AI in Health Systems](#) | Aidoc [Accessed 20 Aug. 2026].



About Indium

We are an AI services company specializing in Agentic AI, Data & Analytics, Application Engineering, and Quality Engineering.

By combining autonomous AI systems, modern engineering, and data intelligence, we build innovative AI solutions that drive measurable business outcomes and long-term value.

With 5,000+ associates worldwide, Indium partners with Fortune 500 companies, Global 2000 enterprises, and leading technology firms in Financial Services, Healthcare, Manufacturing, Retail, and Technology. Our teams operate in North America, India, the UK, Singapore, Australia, and Japan, helping businesses stay ahead in an AI-first world.

USA

Cupertino | Princeton
Toll-free: +1-888-207-5969

INDIA

Chennai | Bengaluru | Mumbai
Hyderabad | Pune
Toll-free: 1800-123-1191

UK

London
Ph: +44 1420 300014

SINGAPORE

Singapore
Ph: +65 6812 7888

www.indium.tech



For Sales Inquiries
sales@indium.tech



For General Inquiries
info@indium.tech

