
The State of Healthcare AI Readiness 2026

Why adoption is outrunning readiness — and what to do about it.

ANKUR JAIN, ESQ., LLB, MBA

FOUNDER — ARTHA CONSULTING LAB · AUTHOR — ARC FRAMEWORK · PODCAST
HOST — THE HUB BRIEF

EXECUTIVE EDITION · 10-MINUTE READ

01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

05 · Agentic Era

SECTION 1 — EXECUTIVE SUMMARY

Organizational readiness — not access to technology — determines whether adoption becomes repeatable performance.

AI use is widespread, yet repeatable organizational and patient value remains hard to scale.

01 Adoption neither creates nor guarantees repeatable value.

P. 03

The adoption and scaling evidence shows that buying and deploying AI is easier than building the capacity to sustain operating improvement and patient benefit. The consequence is a growing portfolio of pilots, rework, and risk without durable improvement in care.

02 Healthcare is buying faster than it can absorb AI.

P. 04

One survey found 85% of US healthcare leaders exploring or adopting generative AI, while a separate comparison places the sector below the cross-industry average for data and AI maturity. The measures are not interchangeable: one captures appetite; the other describes delivery capacity.

03 Governance and workforce authority are hidden constraints on scaling value.

P. 05

Only 21% of companies planning agentic AI within two years report mature agent governance. Weak agentic governance exposes patients to unsafe decisions; weak adoption withholds the benefit the system was meant to create.

04 Readiness can be turned into a funding decision.

P. 06

The ARC Readiness Assessment uses roughly 35 organization-specific questions across Data Strategy, Change Management, and AI Governance. It returns three things a board can act on: the stage the organization can safely absorb, the weakest of the three pillars, and the first evidence-backed fix required before the next AI deployment.

05 Today's foundations determine who shapes the agentic era.

P. 07

Gartner expects task-specific agents in 40% of enterprise applications by the end of 2026. Readiness determines whether an organization negotiates that future or accepts it on a vendor's terms.

THE DEFINITION THIS REPORT WORKS FROM

Readiness is the organizational state in which data, people, and governance can absorb AI without operating performance, patient benefit, trust, or control collapsing.

01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

05 · Agentic Era

FINDING 1

Adoption neither creates nor guarantees repeatable value.

The adoption and scaling evidence shows that buying and deploying AI is easier than building the capacity to sustain operating improvement and patient benefit. The consequence is a growing portfolio of pilots, rework, and risk without durable improvement in care.

Another AI system, platform, or vendor cannot repair an unready organization.

88%

of surveyed organizations report
AI adoption

74%

struggled to achieve and scale
value

SOURCE & WHAT THIS DOES NOT SHOW

McKinsey reports adoption in 88% of surveyed organizations; a separate BCG survey found 74% struggled to achieve and scale value. The samples cannot be combined — one measures adoption, the other measures scaling. Read together, not combined, they establish the leadership question: why does access to similar technology produce different results?

39%

of respondents experimenting
with agents

23%

scaling an agentic system
somewhere in the enterprise

Same McKinsey survey: “The state of AI in 2025: Agents, innovation, and transformation,” November 2025. Experimentation and enterprise-wide scaling are separate measures and are not additive.

01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

05 · Agentic Era

FINDING 2

Healthcare is buying faster than it can absorb AI.

85%

of US healthcare leaders exploring or adopting generative AI

Below

the cross-industry average for data and AI maturity — a directional finding

SOURCE & WHAT THIS DOES NOT SHOW

The appetite signal is McKinsey, “Generative AI in healthcare,” March 2025; the maturity comparison is the World Economic Forum with BCG, “The Future of AI-Enabled Health,” January 2025. The measures are not interchangeable: one captures appetite; the other describes delivery capacity. They are read together, never combined.

When appetite outruns reliable data, prepared roles, and controls, patients experience the gap as delay, inconsistent decisions, and extra coordination. Organizations absorb rework and liability; patients absorb slower access and repeated information requests.

1	Unready data	About 80% of medical data by volume is unstructured. An AI system cannot produce reliable workflow decisions from data the organization cannot locate, connect, or validate.
2	Displaced authority	Pilots clear controlled benchmarks and stall when rollout reaches people whose daily judgment the AI system begins to perform.
3	Immature governance	Static policies age while AI systems, workflows, and regulations change; the organization loses a defensible decision trail.
4	Aggregate measurement	An enterprise average can hide a clinician’s new documentation burden, an intake team’s reconciliation queue, or a patient’s added handoff.



These are organizational conditions, not features of the AI system. Capital then follows what is easiest to buy rather than the constraint most likely to limit organizational performance and patient benefit.

01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

05 · Agentic Era

FINDING 3

Governance and workforce authority are hidden constraints on scaling value.

Weak agentic governance exposes patients to unsafe decisions; weak adoption withholds the benefit the system was meant to create.

21%

of companies planning agentic AI within two years report mature agent governance

SOURCE & WHAT THIS DOES NOT SHOW

Deloitte, “The State of AI in the Enterprise: The Untapped Edge,” January 2026. The 21% figure applies to companies planning agentic deployment within two years. It is a governance-maturity self-report, not a measure of deployed autonomy.

- | | | |
|----------|------------------------------------|---|
| A | Authority moves first | AI can move the first read, interaction check, or prior-authorization call into a system. Defining work leaves while the title, headcount, and liability remain — and training cannot restore judgment that has moved. |
| B | Silence reads as resistance | The loss appears as silence, delay, and workarounds. When leaders label the response as resistance, they prescribe the wrong intervention; adoption weakens and patients receive uneven benefits across teams and shifts. |
| C | Controls must be shown | Leaders must specify where a human stops the system, where guardrails steer it, who answers for the result, and what evidence can be demonstrated. A control that cannot be shown is treated as absent. |



THE FIELD MECHANISM

A professional becomes a confirmer of a machine's judgment at machine speed. Layoffs are counted; lost authority appears on no ledger I have seen.

51% reported at least one negative consequence

30% inaccuracy — the most common consequence

Among respondents at organizations already using AI, McKinsey, “The state of AI in 2025,” November 2025. Without trustworthy data, human oversight, and accountable controls, an organization may scale the error faster than the benefit.

01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

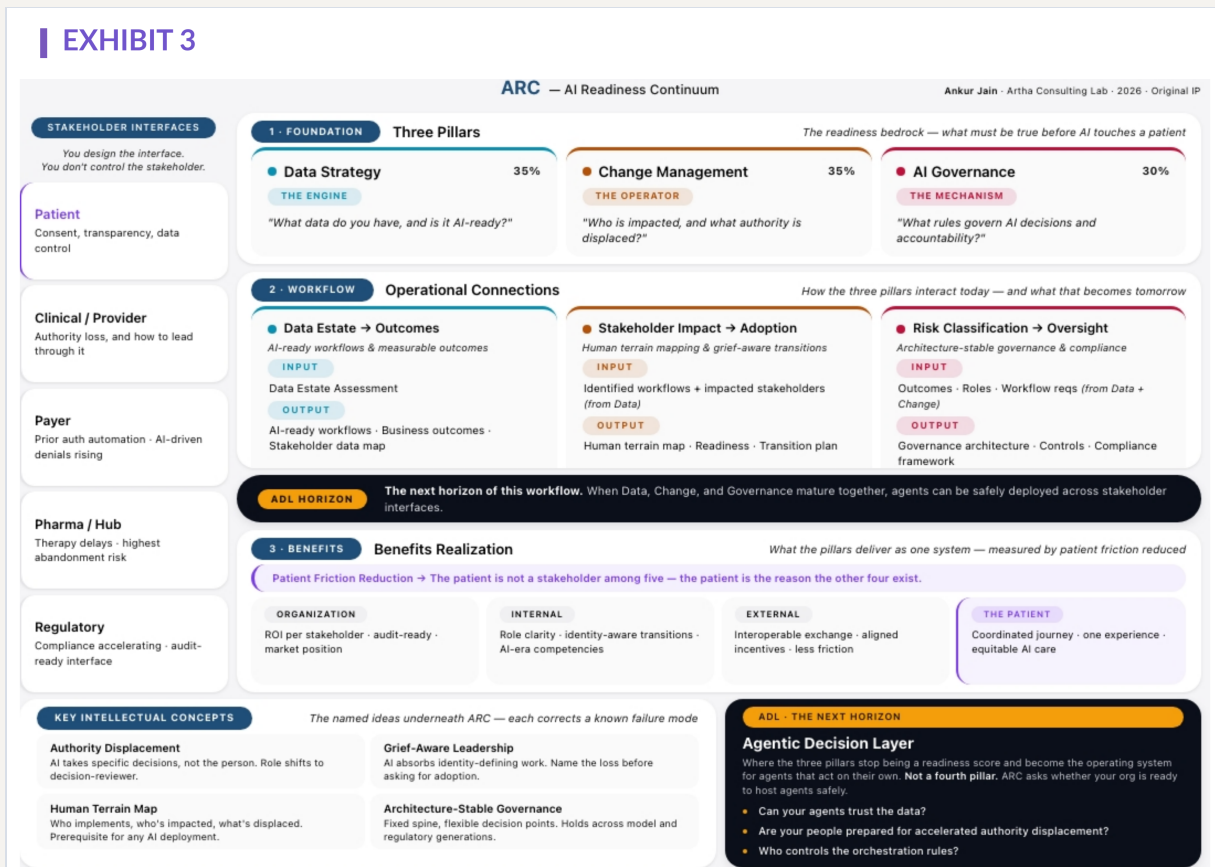
05 · Agentic Era

FINDING 4

Readiness can be turned into a funding decision.

The ARC Readiness Assessment uses roughly 35 organization-specific questions across Data Strategy, Change Management, and AI Governance. It returns three things a board can act on: the stage the organization can safely absorb, the weakest of the three pillars, and the first evidence-backed fix required before the next AI deployment.

The readout does not average the pillars. The five stages are Fragmented, Emerging, Structured, Integrated, and Adaptive, and the weakest-pillar rule sets the overall stage equal to the least mature pillar. Averaging would hide the real absorption limit; bottleneck detection redirects capital to the constraint and names the evidence gate that opens the next rung.



01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

05 · Agentic Era

FINDING 5

Today's foundations determine who shapes the agentic era.

If agents become a coordination layer above current systems, strong foundations will compound operating and patient value; weak foundations will automate fragmentation. Readiness determines whether an organization negotiates that future or accepts it on a vendor's terms.

FORECAST

40%

of enterprise applications
expected to feature task-specific
AI agents by the end of 2026

FORECAST

15%

of day-to-day work decisions
expected to be made
autonomously by 2028

SOURCE & WHAT THIS DOES NOT SHOW

Both figures are Gartner forecasts, not measured deployment counts: "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," August 2025, and the 2028 autonomous-decision projection. Gartner separately expects more than 40% of agentic AI projects to be canceled by the end of 2027, citing cost, unclear value, and inadequate risk controls. Growth inside applications is evidenced; cross-application coordination remains an inference.

The Agentic Decision Layer (ADL) is ARC's name for that proposed destination: a design horizon in which agents coordinate decisions trapped in separate systems. The ADL is not a deployed product. It still depends on trustworthy data, adopted workflows, accountable governance, and a human route to review. Buying agents before those foundations automates fragmentation.

THREE PREDICTIONS FOR 2028

2028 · FOUNDATIONS

Early healthcare operating-layer contracts will be won on data-estate readiness as AI-system capabilities converge. Fragmented organizations will negotiate as customers of someone else's platform.

2028 · VENDOR ECONOMICS

Outcome-linked structures will begin to displace per-seat pricing in new healthcare AI contracts as agents reduce the human licenses that pricing model was built to meter.

2028 · GOVERNANCE

Organizations automating day-to-day decisions without agent-ready governance will reproduce today's harms at machine speed.

01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

05 · Agentic Era

SECTIONS 4, 6 & 8 — DECISIONS FOR LEADERS

Measure the pillar leadership is least willing to examine.

Then decide how long staff and patients should keep paying for the uncertainty. Three decision boxes from the full report, one per audience.

BEFORE APPROVING A USE CASE

Health System

Identify the workflow, the limiting pillar, the people whose decision rights change, and the patient measure. Without those four answers, the proposal is still a pilot bet.

Biopharma / Hub

A hub unable to recognize one patient across its own systems lacks a foundation for safe automation. Fix identity and handoffs before the patient pays for another layer of fragmentation.

Payer

Deep automation without an accountable owner, evidence, and an appeal path will look like denial optimization. Use the patient journey test to prove the coverage workflow is improving.

BEFORE THE NEXT BOARD DECK

Health System

Put a stage, bottleneck, first fix, and patient measure in the next board deck. If no one can name the weakest pillar, that uncertainty is itself a finding.

Biopharma / Hub

Assess identity and journey data before buying the next automation. An AI system cannot recognize a patient whom the data estate cannot connect.

Payer

Deep automation may flatter Data Strategy. Measure decision rights, adoption, appeals, and overrides just as rigorously; member impact appears where those controls fail.

MONDAY-MORNING ACTIONS

Health System

Name the next two workflows and assign a human-terrain map before vendor demonstrations. No owner by Friday is evidence that Change Management is unready.

Biopharma / Hub

Request one patient journey from prescription through therapy. If producing it requires a meeting, prioritize identity resolution and journey mapping.

Payer

Inspect the most automated workflow's governance record. If it lacks an accountable decision owner, evidence trail, appeal, and override, make that gap the year's priority.

Across sectors, the agenda should reduce the integration work carried by patients.

01 · Adoption

02 · Capacity

03 · Governance

04 · Assessment

05 · Agentic Era

SECTION 8 — THE ACTION PLAN

Each quarter creates the prerequisite for the next.

Measure the weakest pillar, repair it, deploy only what the evidence supports, then close the feedback loop.

MONTHS 1–3

Establish the readiness baseline

Measure the stage, name the bottleneck, and make changed authority visible before another deployment decision.

MONTHS 4–6

Repair the limiting pillar

Pause work that depends on the bottleneck until evidence shows that the constraint changed.

MONTHS 7–9

Deploy inside the architecture

Deploy only after benefits, burdens, decision rights, monitoring, and patient safeguards are explicit.

MONTHS 10–12

Close the loop and reassess

Turn a year of projects into a standing operating rhythm. The board should see a readiness trend and an evidence gate — not another demo count.


EXHIBIT 11

THE STATE OF HEALTHCARE AI READINESS 2026

A 12-month sequence with a decision gate every quarter

Each quarter produces the evidence required to unlock the next.

THE LOOP RETURNS LEADERS TO THE NEXT READINESS MEASUREMENT



THE SEQUENCE IS FIXED. THE STARTING HYPOTHESIS VARIES.

Test the hypothesis against the readiness baseline.

HEALTH SYSTEM / CIO-CMO	LINKED GATE: BASELINE ACCEPTED	BIOPHARMA / HUB	LINKED GATE: BOTTLENECK PROOF	PAYER	LINKED GATE: GO-LIVE CONDITIONS
STARTING HYPOTHESIS Human terrain + authority QUESTION TO TEST Who owns the two workflows likely to deploy—and whose judgment moves? PROOF REQUIRED Named owner + completed human terrain map		STARTING HYPOTHESIS Identity + journey data QUESTION TO TEST Can one patient be followed from prescription through therapy? PROOF REQUIRED Shared identity + end-to-end journey evidence		STARTING HYPOTHESIS Decision rights + audit evidence QUESTION TO TEST Can an automated denial be traced to evidence, authority, and override? PROOF REQUIRED Named accountability + reviewable AI-system and human audit trail	

Starting hypotheses only. If the baseline names a different weakest pillar, follow the evidence—not the sector label.

Artha Consulting Lab · 2026

Exhibit 11 · The 12-month readiness sequence and its decision gates

Exhibit · from the full report

The four decision gates stay fixed; the starting hypothesis changes by organization and must yield to baseline evidence. Judge spending by whether it removes the constraint, protects the patient, and clears the next evidence gate.

NEXT STEP

This is the 10-page edition. The full report goes further.

The full 51-page report carries what this edition could not: all 11 exhibits, the benchmark of published readiness frameworks against three completeness tests, the five-stage measurement logic, the composite readiness profiles, and the complete method and sources. Download it free at arthaconsultinglab.com.

arthaconsultinglab.com — download the full report

Downloading also brings the ARC white paper first when it releases, and an invitation to a complimentary 20-minute readiness conversation. For the framework brief, message me on LinkedIn or write to ankur.jain@arthaconsultinglab.com.

About the Author

Ankur Jain, Esq., LLB, MBA

Founder — Artha Consulting Lab · Author — ARC Framework · Podcast Host — The Hub Brief

Ankur Jain is founder of Artha Consulting Lab and author of the AI Readiness Continuum. For more than a decade, he has worked across pharma, hubs, specialty pharmacies, payers, and health systems on post-deployment data integration, workflow adoption, decision accountability, governance, and value realization. He hosts The Hub Brief and writes on the human side of healthcare AI.

Artha Consulting Lab helps healthcare organizations measure readiness, identify the weakest pillar, and strengthen the constraint holding back organizational and patient value. It also builds governance designed to survive the next AI system and regulation.

arthaconsultinglab.com · linkedin.com/in/ankurjaincons · The Hub Brief on YouTube, Substack, and Apple Podcasts

About This Edition

This Executive Edition condenses the inaugural full report. Quantitative claims come from published sources with their population qualifiers attached; field observations are labeled as such and are not presented as survey evidence; 2028 statements remain predictions, not measured outcomes.

Suggested citation: Jain, Ankur. *The State of Healthcare AI Readiness 2026: Why adoption is outrunning readiness — and what to do about it*. Inaugural edition. Artha Consulting Lab, 2026.