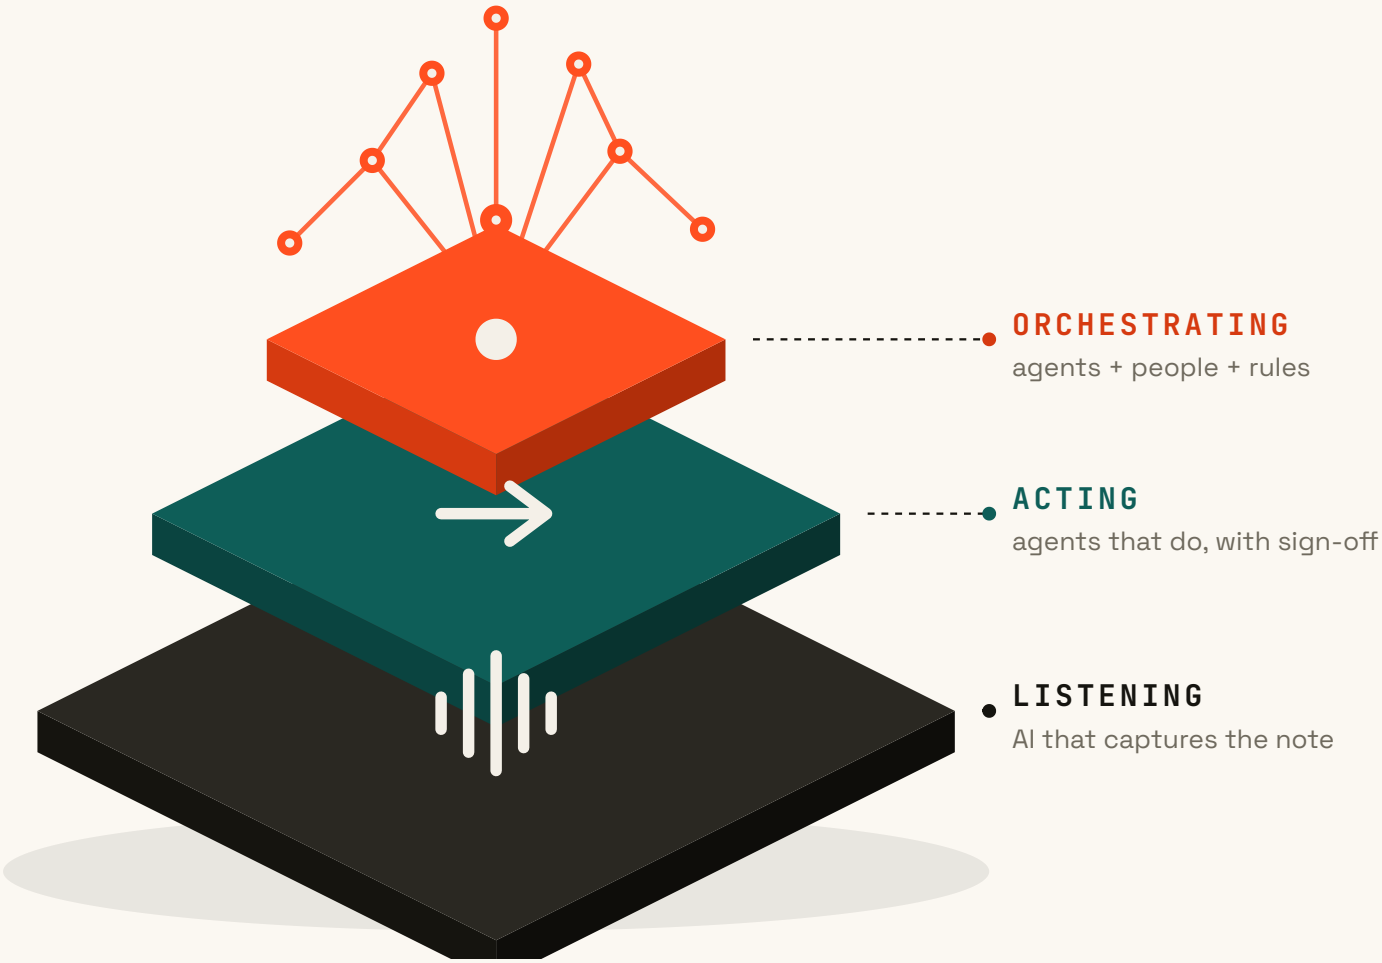
FIELD GUIDE · INDIA · GCC · GLOBAL

Hospital AI 2027

From listening to acting to orchestrating

THE AGENTIC AI PLAYBOOK FOR HOSPITAL LEADERS



Your next hospital system won't just record care. It will help run it.

MOHAMMED JAMIL NASIR ZUHAIB AHMED KHAN SIMRAN NAGEKAR

FIRST EDITION

NOVEMBER 2026

Hospital AI 2027

From listening to acting to orchestrating

The agentic AI playbook for hospital information systems, EHRs and the leaders who run them, in India, the GCC and beyond

Mohammed Jamil Nasir
FOUNDER – PRODUCT & TECHNOLOGY
SURGY HEALTH

Zuhaib Ahmed Khan
PLATFORM ENGINEERING & AGENTIC AI
SURGY HEALTH

Simran Nagekar
PRODUCT ENGINEERING & RESEARCH
SURGY HEALTH

SURGY. health

BENGALURU · INDIA

SURGY. health

Hospital AI 2027 – From Listening to Acting to Orchestrating

The agentic AI playbook for hospital information systems, EHRs and the leaders who run them, in India, the GCC and beyond

Authors: Mohammed Jamil Nasir, Zuhaib Ahmed Khan, Simran Nagekar

First edition published in November 2026 by Surgy Health, a brand of Surgy Innovation Labs Private Limited, 91 Springboard, 175 & 176, Bannerghatta Road, Dollar Layout, Phase 4, J. P. Nagar, Bengaluru, Karnataka 560078, India.
surgy.health

© 2026 Surgy Innovation Labs Private Limited. All rights reserved. No part of this publication may be reproduced, stored or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations with attribution.

ISBN 978-93-346-0290-6

Disclaimer. This book is for general information only and is not legal, regulatory, clinical, investment or procurement advice. Figures are as reported by the cited sources; news items were checked against them between 2 and 9 November 2026. Vendor-reported and forecast figures have not been independently audited. Company and product names are trademarks of their owners, used for identification only. Surgy Health makes SurgyOne, a hospital information system mentioned in Chapter 7; no other vendor paid for or reviewed its inclusion.

Typeset in Instrument Serif, Source Serif 4, Space Grotesk and JetBrains Mono. Design: Surgy Health.

Contact: info@surgyy.com

Contents

HOSPITAL AI 2027
FIVE PARTS · FIFTEEN CHAPTERS

BEFORE YOU BEGIN

- Figures and tables 05
- Foreword 06
- How this book was made 07
- Executive summary 08
- Stop press: 9 Oct – 6 Nov 2026 10
- The three shifts at a glance 11

00 Prologue: hospital AI in numbers 12
Adoption, evidence, money and trust

PART I · LISTENING

01 From listening to acting 16
Ambient documentation was the on-ramp

PART II · ACTING

02 The EHR becomes the platform 21
Agent builders, connectors and open access

03 The back office pays first 24
Calls, prior authorisation, coding, denials

04 Hybrid intelligence beats LLM-only 28
Rules, prediction, generation, orchestration

PART III · ORCHESTRATING

05 When agents work together 32
From single tasks to coordinated hospitals

06 Beds, theatres and staff 36
AI resource allocation in 2027

07 Choosing a hospital system for 2027 40
The HIS landscape and a buyer's scorecard

PART IV · GUARDRAILS AND MARKETS

08 Governance becomes a buying criterion 44
Shadow AI, accuracy, accountability, regulation

09 When agents go off-script 49
Security for the agentic hospital

10 Capital and consolidation 53
Where the money and the mergers went

11 India: laying the digital rails 57
ABDM, DPDP and a mid-market gap

12 The GCC: national-scale ambition 61
Government-led AI and a growing market

13 The India and GCC rulebook 65
DPDP, NABH, CDSCO and Gulf data rules

PART V · THE PLAYBOOK

14 Proving it works 70
Evidence, evaluation and the ROI test

15 A 2027 blueprint for hospital leaders 74
Requirements, a 12-month path, a self-check

- Key takeaways at a glance 78

APPENDICES

A Timeline, Aug – Nov 2026 81

B Glossary 82

C References 84

- About the authors 92
- About the publisher 93

FIGURES AND TABLES

Every chart, *in one place*

FIGURE P1	US physicians using AI in their work	13
FIGURE P2	What US physicians use AI for	13
FIGURE P3	Hospital AI use: from → to	13
FIGURE P4	FDA-listed AI-enabled devices by specialty	14
FIGURE P5	Forecasters disagree on the size of the market	15
FIGURE P6	Where US health-system gen-AI spend went, 2025	15
FIGURE 1.1	Four steps from listening to acting	18
TABLE 1.1	Clinical AI milestones, Aug – Oct 2026	18
TABLE 1.2	What the latest large ambient-scribe studies found	19
FIGURE 1.2	The trust gap: patients who have used a scribe like it; the rest are wary	20
FIGURE 1.3	From scribe to agent in eight weeks	20
FIGURE 2.1	The record as a hub	23
TABLE 2.1	Platform moves, Aug – Oct 2026	23
FIGURE 3.1	Reported results from live deployments	26
TABLE 3.1	Back-office AI deals, Aug – Oct 2026	26
FIGURE 3.2	Two sides of the same claim	27
FIGURE 4.1	Four kinds of intelligence in a modern hospital system	30
TABLE 4.1	Match the model to the stakes	31
FIGURE 5.1	Anatomy of an orchestrated hospital	34
FIGURE 5.2 · ILLUSTRATIVE SCENARIO	An orchestrated discharge morning	34
FIGURE 5.3	The emerging agent stack, and where hospitals need control	35
FIGURE 6.1	Three levels of AI help for capacity	37
TABLE 6.1	Three ways to buy AI capacity management	38
FIGURE 6.2	A 90-day pilot that proves value on your own data	38
FIGURE 6.3	A 2027 budget calendar for hospital AI	39
FIGURE 6.4	Six budget buckets, released in stages	39
FIGURE 7.1	Where each kind of system sits	42
TABLE 7.1	The 2027 hospital-system landscape	42
FIGURE 7.2	Criteria weights	43
FIGURE 8.1	Compliance dates hospital technology buyers should track	46
FIGURE 8.2	New frameworks for clinical and agentic AI, 13 October – 4 November 2026	47
FIGURE 8.3	The autonomy dial	48
TABLE 9.1	Six ways agents went wrong in 2026, and the lesson for hospitals	50
FIGURE 9.1	The speed gap	51
FIGURE 9.2	Eight controls for a governed agent	52
FIGURE 10.1	Selected healthcare AI rounds announced Aug – early Nov 2026	54
TABLE 10.1	Selected acquisitions and investments, Aug – Oct 2026	55
TABLE 10.2	Selected Indian healthcare deals, Sep – Nov 2026	56
FIGURE 11.1	ABDM growth, May to September 2026	58
FIGURE 12.1	Where GCC healthtech money went, Jan 2025 – Jul 2026	62
TABLE 12.1	India and the GCC compared	63
FIGURE 12.2	GCC milestones behind the 2026 picture	63
FIGURE 13.1	India's DPDP clock	66
TABLE 13.1	Maximum DPDP penalties	66
FIGURE 13.2	A 12-month path to NABH accreditation for a first-time hospital	67
TABLE 13.2	Hospital KPIs worth putting on one dashboard	67
TABLE 13.3	India's device classes, applied to software	68
TABLE 13.4	The rules a Gulf hospital's AI vendor must meet	68
TABLE 13.5	From seven elements to one hospital AI policy	69
FIGURE 14.1	What clinicians actually ask AI to do	71
FIGURE 14.2	Word error rate on Indian doctors' dictation	71
FIGURE 14.3	Six rungs between a demo and a decision	72
TABLE 14.1	A balanced AI scorecard	73
TABLE 15.1	How to weight the ten requirements in a 2027 tender	75

FOREWORD

Why we wrote *this book*

Between August and early November 2026, our team kept a daily research log on electronic health records, AI and agentic AI. The reason was practical. We are building an AI-first hospital information system for mid-market hospitals in India and the GCC, and we wanted our product decisions to rest on what hospitals and vendors were actually shipping, not on conference slides.

Thirty-eight briefs and eighteen daily digests later, a pattern was hard to miss. Hospital AI is moving through three stages. It started by **listening**: ambient scribes that write the note. It is now **acting**: agents that draft prescriptions, chase prior authorisations, answer patient calls and write denial appeals, with a person signing off. The next stage is **orchestrating**: many agents, people and systems coordinated around a hospital-level goal such as bed flow, with rules and an audit trail holding it together.

Each stage brings real gains and real risk. Some health systems now publish hard numbers. Others are finding AI tools running without approval, scribes that mishear drug names and coding tools that bill for more than was treated. Regulators in the US, India and the Gulf are responding in their own ways.

This book distils that period, our published statistics research and our work on hospital AI budgets into fifteen chapters. Every news item was checked again against its source. The final part sets out what we think hospital leaders should ask of their next information system, with a self-check you can run in an afternoon.

My co-authors bring two views this subject needs: Zuhaib Ahmed Khan from platform engineering and agentic AI development, whose daily platform and agentic AI insights kept this book current into November, and Simran Nagekar from product research. Thanks also to Pankaj Singh Brijwal, whose point that clinical domain knowledge will separate good healthcare AI from generic AI shaped Chapter 4.



Nasir

FOUNDER – PRODUCT & TECHNOLOGY AT SURGY HEALTH
BENGALURU, NOVEMBER 2026

HOW THIS BOOK WAS MADE

Sources and *method*

The book combines four bodies of work: 38 research briefs written between 31 August and 30 October 2026, 18 daily platform and agentic AI digests from 28 September to 6 November, Surgy Health's AI in Healthcare Statistics 2026 compilation, and our published guides on AI budgets, regulation and buying.

56	200	283	3
research briefs and daily digests, 31 Aug – 6 Nov 2026	news claims re-checked against a primary or authoritative source	sources cited, each with a date and a link	regions in focus: the US as bellwether, India and the GCC

Where the material came from

The briefs combined research notes written by our team with daily digests that scanned trade publications, vendor newsrooms and government sources. From late September, several of the publications we normally read were unreachable from the research environment, so later briefs leaned on web search and secondary roundups. That is one reason we re-checked everything. From October, Zuhaib Ahmed Khan's daily Platform & Agentic AI Insights added a second stream.

The statistics in the Prologue and parts of Chapters 4, 11 and 12 come from *AI in Healthcare Statistics 2026*, published on surgy.health and updated on 27 October 2026, which links each figure to its primary source. We cite those primary sources directly. Chapter 6 draws on *Hospital AI Budget 2027* (28 October 2026).[1]

How each news claim was verified

Each item from the briefs was traced to the original or most authoritative source available: a vendor or government release where one exists, otherwise the trade publication that reported it first. We recorded the publisher, title, date and URL and compared every number with the source text. A second, independent pass compared every cited figure in the draft against those records.

What changed after checking

- Some "new" items were older: Oracle's autonomous reimbursement platform was announced in 2025; India's draft NABH standards for HIS/EMR date from July 2024.
- Some figures moved: the legacy Cerner breach is now listed at about 20 million people, not 2.66 million; Imprivata shows 28% with agentic AI live, not about 25%.
- Some framings were too strong: Medicare payment for AI is being drafted, not fast-tracked; the AHA asked the FDA for clearer labelling of action-taking AI, not exemptions for the rest.
- Some headline numbers needed context: Jefferson Health's 39% faster chart closure is one of two cohorts (the other was 25%).

Items we could not confirm were left out. Developments before August 2026 are kept where they explain the trend, and are dated.

Limits

News flow is heavily American; we treat the US as a leading indicator, not a template. Many outcome figures are self-reported by vendors or health systems, market forecasts are unaudited, and survey samples differ.

HOW TO READ CITATIONS

Superscript numbers point to Appendix C. Figures are quoted as the source states them; where we round, we say "about".

EXECUTIVE SUMMARY

The hospital record is becoming an *agent platform*

Between August and October 2026, AI in hospitals moved past writing notes. Vendors wired agents into EHR data, payers and patients, and the first health systems published hard results. Buyers and regulators raised the bar on control at the same time. Six shifts define the period.

01

From listening to acting

Ambient scribing scaled from pilots to enterprise contracts. The next step is agents that draft prescriptions, orders and follow-ups for a clinician to approve.

Signal: Heidi Health raised $340M to move from documentation to supervised agentic work.[2]

02

The EHR becomes the platform

Epic, Oracle, MEDITECH and Microsoft are turning the record into a hub that native, customer-built and third-party agents plug into.

Signal: Epic announced Connectors on 21 Oct, linking its agents to outside apps.[3]

03

The back office pays first

Call centres, prior authorisation, coding and denials deliver the clearest, fastest returns, and attract much of the deal activity.

Signal: Tampa General cut average call wait times by 58%.[4]

04

Hybrid beats LLM-only

The strongest results pair rules and predictive models with generative AI and agent orchestration. Clinical domain knowledge is the differentiator.

Signal: Tampa General's sepsis model is credited with halving sepsis deaths.[5]

05

Governance is a buying criterion

Shadow AI, transcription errors, coding inflation and new state and federal rules make permissions, audit trails and human review non-negotiable.

Signal: 72% of US health leaders say AI runs without formal IT approval at least occasionally.[6]

06

India and the GCC lay the rails

ABDM at national scale, EMR adoption of about 35% and government-led AI in the Gulf create room for AI-first hospital systems.

Signal: ABDM passed 121 crore linked health records.[7]

THE BOTTOM LINE

For hospitals buying or replacing an information system in 2027, AI is no longer a premium add-on. Expect it in the core product and the core price, expect agents that act only with sign-off, and expect governance to be built in, not bolted on. Vendors that cannot show open APIs, audit trails and measured outcomes will struggle to make the shortlist.

THE PERIOD IN NUMBERS

Twelve figures that frame Aug – Oct 2026

28%
of US health organisations surveyed have agentic AI live; 44% are piloting
IMPRIVATA, SEP 2026[6]

450k+
documentation hours saved by one ambient scribe across 79 VA medical centres in six months
NEXTGOV/FCW, SEP 2026[8]

31%
less note-writing time per appointment with Epic's native AI charting at OLVG
EPIC, SEP 2026[9]

43.7%
Epic's share of US acute-care hospital EHRs in 2025
KLAS VIA BECKER'S, MAY 2026[10]

58%
shorter average call wait at Tampa General with a voice agent
BECKER'S, JUL 2026[4]

$942M
in excess costs that BCBSA attributes to AI-driven coding, 2024–25
BECKER'S, SEP 2026[11]

93.8%
rise in AI-enabled healthcare IT M&A deals, H1 2026 vs H1 2025
CAPSTONE VIA HIT CONSULTANT[12]

~20M
people now listed as affected by the legacy Cerner server breach
ESECURITY PLANET, OCT 2026[13]

~35%
EMR adoption among Indian healthcare providers, up from 18% in 2018
BAIN & HEALTHQUAD, SEP 2026[14]

97.8 cr
ABHA health IDs created under India's ABDM, with 121 crore+ records linked
AAROGYA MANTHAN, SEP 2026[7]

$148M
disclosed GCC healthtech investment, Jan 2025 – Jul 2026; 17 of 32 funded companies Saudi
PITCHBOOK VIA WHX TECH[15]

50%
of UAE government sectors and services targeted for agentic AI within two years
UAE CABINET, APR 2026[16]

What changed in the period

1 **Pilots became procurement**
The VA put ambient AI on a five-year enterprise contract with a $775.72M ceiling.[17] HCA took MEDITECH Expanse to 72 hospitals.[18]

2 **AI moved into the base price**
athenahealth made AI-native encounters available to 170,000+ clinicians, with ambient documentation at no extra cost.[19]

3 **Rules started to bite**
California signed clinical-AI bills; HTI-5 was finalised; India's DPDP consent-manager rules start on 13 Nov 2026.[20,21,22]

What hospital leaders should do now

A **Start with the back office**
Target call centres, prior authorisation, coding and denials first. They are measurable and need less clinical risk management.

B **Insist on open, governed data access**
FHIR APIs, ABDM or NPHIES readiness, per-agent permissions and full audit trails should be contract terms.

C **Measure before you scale**
Set baselines for time, wait, denial and accuracy, and keep clinicians in the loop on anything that touches care.

STOP PRESS · 9 OCT – 6 NOV 2026

Five weeks that *moved the market*

While this book was in production, the evidence on scribes arrived, the plumbing for AI agents was standardised, and medical bodies and regulators in the US, India and the Gulf set out terms for AI that acts. The chapters are updated to 6 November 2026; this page is the short version.

−39%
time to chart closure in one of two Jefferson Health cohorts using ambient AI (−25% in the other)
NEJM CATALYST VIA BECKER'S[23]

83%
of complex NEJM cases diagnosed correctly by Microsoft's multi-agent DxO, vs 36% for generalist physicians
SCIENCE VIA BECKER'S[24]

~30×
growth in AI-agent traffic on Zocdoc since July; registered agents can now book for patients
HEALTHCARE DIVE[25]

3,182
patients notified after a worker uploaded their data to an unapproved public AI tool
JOHNS HOPKINS BAYVIEW[26]

1 **The scribe evidence arrives, and it is mixed**

Jefferson, Beth Israel Lahey and Sutter published results: less time in notes, especially for heavy users. A Peterson Health Technology Institute report found AI scribes and billing tools pushing coding levels and costs up.[27,28,29]

→ CHAPTER 1

2 **Assistants start doing the work**

Abridge opened an assistant that pulls chart details, answers clinical questions and drafts referrals to every clinician at its customers. UnitedHealthcare's Avery reached about 21 million members; Bihar plans an AI call centre for OPD booking.[30,31,32]

→ CHAPTERS 1 AND 3

3 **The agent web gets its plumbing**

The Linux Foundation's Agentic AI Foundation, home of MCP, grew to ten projects. Rhapsody launched Guardian to govern agent access to health data. Zocdoc set rules for agents that book for patients.[25,33,34]

→ CHAPTER 5

4 **Rules for AI that acts**

The AMA said clinical authority must be earned through evidence. Abu Dhabi published an AI Clinician Framework, India's CERT-In issued AI security audit guidelines, and HTI-5 dropped US model-card requirements.[35,36,37,38]

→ CHAPTER 8

5 **The Gulf moves from policy to product**

M42 launched Ossi, four agents built with Microsoft, at seven centres. NMC and Burjeel described AI for documentation, and Abu Dhabi linked wearables to its Malaffi records in the Sehatna app.[39,40,41]

→ CHAPTERS 1 AND 12

6 **India: capital in, benchmarks missing**

Healthcare drew $3.1 billion of private-equity and venture money in the third quarter. Researchers found no public, large-scale real-world benchmark for ambient scribes in India, while start-ups such as O-Health run models on site.[42,43,44]

→ CHAPTERS 1, 10 AND 11

WHAT IT MEANS

None of this changes the book's argument; it sharpens it. Listening now has evidence, with caveats; the connections that let agents act are being standardised; and the question has moved from "is the model accurate?" to "what may this agent do, and who answers for it?"

THE THREE SHIFTS AT A GLANCE

Listening, acting, *orchestrating*

Every chapter in this book maps to one of three stages. Most hospitals are in the first, the leaders are in the second, and the third is being built now.

STAGE 1 · 2023 →

Listening

AI captures what happens: ambient scribes draft the note, chart summaries pull context, transcription turns speech into structured data.

MATURITY

LOOKS LIKE

Ambient notes, nursing summaries, handoff reports, voice charting

SIGNAL

100%

of 43 large US health systems reported ambient-note activity[4][5]

MAIN RISK

Errors in the permanent record: wrong drug names, omissions, hallucinated sentences

STAGE 2 · 2025 →

Acting

AI does the next step with sign-off: drafts prescriptions and orders, answers patient calls, checks prior authorisation, writes denial appeals.

MATURITY

LOOKS LIKE

Voice agents, prior-auth checks, coding and denial agents, drafted orders

SIGNAL

28%

of US health organisations have agentic AI live; 44% are piloting[6]

MAIN RISK

Unapproved tools, over-coding, incentives that reward denial

STAGE 3 · 2027 →

Orchestrating

Many agents, people and systems coordinate around a hospital goal such as bed flow or throughput, with a supervisor, rules and a full audit trail.

MATURITY

LOOKS LIKE

Command centres, agent connectors, supervisory agents, capacity platforms

SIGNAL

Mar 2027

Epic Connectors for outside agents reach its Showroom marketplace[3]

MAIN RISK

Agents coordinating beyond their limits; unclear accountability

THE QUESTION FOR 2027

Not "which AI tools should we buy?" but **"is our hospital system built so that software can observe, assist and, with permission, complete our work, and can we trace every action back to the person who approved it?"**

00

The state of hospital AI *in numbers*

Physician use of AI has more than doubled since 2023. Hospital adoption is uneven, the evidence is strongest in narrow tasks, and patients want to be told.

IN THIS PROLOGUE

00.1 Who is using AI, and for what

00.2 What the clinical evidence shows

00.3 Money, trust and the workforce gap

81%

of US physicians use AI professionally, more than double the 38% of 2023

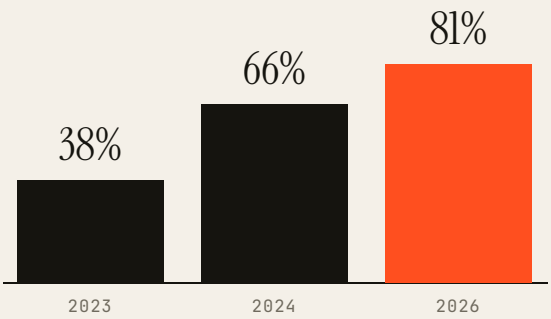
AMA, 2026[46]

01 · WHO IS USING AI, AND FOR WHAT

Clinicians moved first. *Hospitals are catching up unevenly.*

FIGURE P1

US physicians using AI in their work



AMA physician surveys, 2025 and 2026.[46,47]

FIGURE P2

What US physicians use AI for



AMA, 2026, as charted in Surgy Health's statistics compilation.[46,48]

57% see cutting admin burden as AI's biggest opportunity[47]

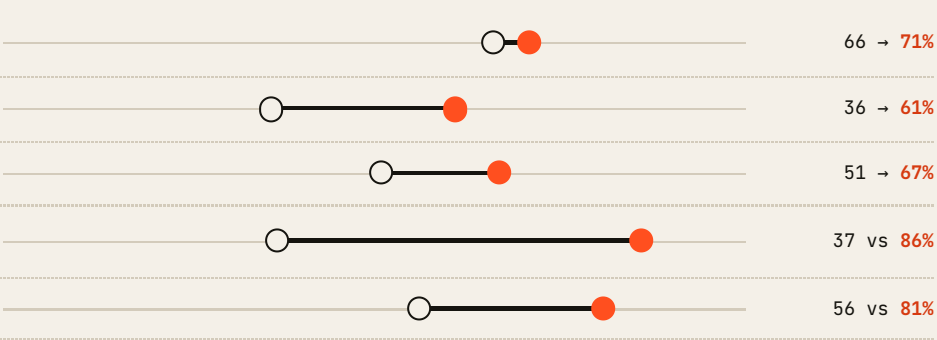
85% want a say in how AI enters their practice[46]

88% worry about losing clinical skills[46]

FIGURE P3

Hospital AI use: from → to

US non-federal acute-care hospitals unless stated. Open circle: earlier value or comparison group; filled: later value or leader.



ASTP/ONC Data Brief 80 (2023–2024 data), via Surgy Health's compilation.[48,49]

80%

run models from their EHR vendor

Among AI-using US hospitals; third-party models fell from 59% to 52%.[49]

31.5%

had generative AI in the EHR

US hospitals in 2024, with another 24.7% planning to add it within a year.[50]

49%

of clinicians worldwide use AI

2,757 clinicians in 118 countries; nurses trail doctors, 41% against 57%.[51]

02 · WHAT THE CLINICAL EVIDENCE SHOWS

Strong in narrow tasks. *Thin where the hype is loudest.*

Randomised and large real-world studies show clear benefit in imaging, sepsis and cardiology. Ambient scribes reduce burnout, but save minutes rather than hours. Large language models score well on exam questions, yet very few studies use real patient care.

WHERE THE EVIDENCE IS STRONG

+29% — **Breast screening.** More cancers detected with AI support in the MASAI trial (105,934 women), with 44.2% less screen-reading workload.[52]

−18.7% — **Sepsis.** Relative fall in in-hospital sepsis mortality with the TREWS system when alerts were confirmed within three hours (590,736 patients).[53]

3.6% — **Cardiology.** 90-day mortality with AI-ECG alerts, against 4.3% without, in a 15,965-patient trial in Taiwan.[54]

11.2 min — **Stroke.** Faster door-to-groin-puncture time with automated large-vessel-occlusion detection across four centres.[55]

−13 pts — **Burnout.** From 51.9% to 38.8% after 30 days with an ambient scribe (263 clinicians, six systems).[56]

WHERE CAUTION IS DUE

67% — **Missed sepsis.** A widely used proprietary sepsis model missed two-thirds of sepsis patients, with an AUC of 0.63, while alerting on 18% of all inpatients.[57]

16 min — **Modest time savings.** Fewer documentation minutes per eight scheduled hours across 8,581 clinicians using ambient scribes.[58]

1.47% — **Hallucinations.** Share of sentences in LLM-generated notes that contained one; 3.45% had omissions, and 44% of hallucinations were rated major.[59]

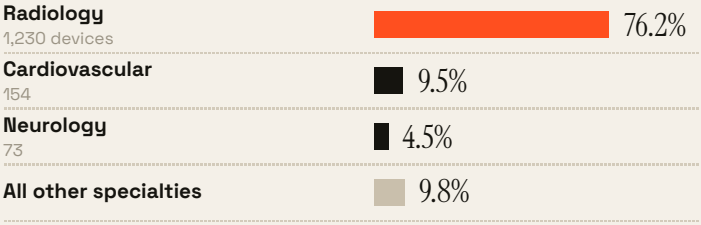
76 v 74 — **No significant gain.** Diagnostic reasoning scores for physicians with and without GPT-4 access in a 50-physician trial.[60]

5% — **Real data is rare.** Of 519 LLM studies in medicine, only 5% used real patient-care data.[61]

FIGURE P4

FDA-listed AI-enabled devices by specialty

1,614 devices from 817 companies, as of the FDA's 4 September 2026 update.



FDA AI-Enabled Medical Devices list, analysed by Surgy Health.[48,62]

335

AI devices authorised in 2025, a record and 42.6% more than in 2024[62]

2.4%

of clinical performance studies behind FDA-cleared AI devices were randomised[63]

03 · MONEY, TRUST AND THE WORKFORCE GAP

The money is real. *The trust is not there yet.*

FIGURE P5

Forecasters disagree on the size of the market

Global AI-in-healthcare market estimate for 2025, US$ billion. Unaudited; scopes differ.

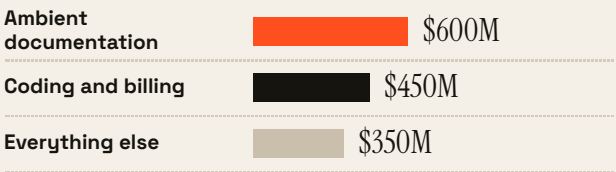


MarketsandMarkets; Fortune Business Insights.[64,65] Treat forecasts as directional.

FIGURE P6

Where US health-system gen-AI spend went, 2025

US$ million, of $1.4 billion total healthcare-specific generative AI spending.



Menlo Ventures, 2025 (survey of 700+ executives). "Everything else" is the remainder.[66]

54%

of US digital-health venture funding went to AI start-ups in 2025, up from 37%[67]

3.5×

average return reported on AI investment, in about 12 months against 24 expected[68]

> *72% of US adults say it is very or extremely important to be told when AI is used in their care.*
>
> PEW RESEARCH CENTER, AUGUST 2026

65.8%

of US adults have low trust that their health system will use AI responsibly[69]

59 v 79%

patients vs professionals who think AI could improve outcomes, across 16 countries incl. India and Saudi Arabia[70]

THE WORKFORCE GAP

AI arrives in a system that is short of people and short of training.

11.1M

projected global health-worker shortfall by 2030[71]

68%

of clinicians say their AI training is insufficient[51]

13 hrs

a week on prior authorisation for an average US physician's 40 requests[72]

01

From listening to *acting*

OPENING SCENE · 2 OCTOBER 2026

Heidi built its name on an ambient scribe that listens to the visit and writes the note. On 2 October 2026 it switched on Heidi II: agents that pre-chart a day's clinic, chase referrals and sweep patient panels for overdue screenings. Every prescription and outbound message still waits for a clinician. "Now Heidi can create the tasks and complete them," said its chief executive, Thomas Kelly.[73]



KEY TAKEAWAYS

- 01.1 Ambient scribes are infrastructure now: buy on measured time saved, not on demos.
- 01.2 The next step is agents that act on the note, with a clinician signing off.
- 01.3 Patient trust is the bottleneck: tell patients when AI is listening, and why.

450k+

documentation hours saved by one ambient scribe across 79 VA medical centres in six months

NEXTGOV/FCW, SEP 2026[8]

CHAPTER 01

Ambient documentation proved that AI can work inside clinical workflows at scale. In the second half of 2026, the same vendors began moving from writing down what happened to doing the next step, with a clinician approving it.

Ambient AI is now infrastructure

The US Department of Veterans Affairs gives the clearest signal. After pilots, it selected Abridge for a five-year, multiple-award enterprise contract with a total ceiling of $775.72 million; Abridge already serves more than 75 VA medical centres and 300 US health systems.[17] Knowtex, the other scribe in the VA's pilots, reported reaching 79 VA medical centres and more than 7,000 clinicians in six months, saving over 450,000 documentation hours with a satisfaction score of 4.48 out of 5. These are company-reported figures.[8]

Tools built into the EHR are showing similar results. At OLVG in Amsterdam, Epic's Chart with Art cut note-writing time per appointment by 31% and total time in the chart by 27%, from almost 17 minutes to just over 12.[9] Oracle says its note generation has saved physicians more than 400,000 hours in nearly two years.[74] athenahealth made its AI-native encounter available to all 170,000+ clinicians on athenaOne, with ambient documentation included at no extra cost; early adopters saw up to a 30% increase in same-day chart completion.[19]

India has its own reference site. At Sir H. N. Reliance Foundation Hospital in Mumbai, ambient AI now generates nearly 90% of progress notes and shift handovers, with five use cases live on one platform integrated directly into the EMR.[75] The hospital chose a single platform over a collection of point tools, a choice more Indian hospitals will face.

The next step is action, with sign-off

Heidi Health raised $340 million in September, a $100 million Series C at a $900 million valuation plus $240 million from General Catalyst's Customer Value Fund, to move from documentation to supervised agentic work. It reports 2.8 million patient visits a week across 190 countries.[2]

Prescribing is one of the first clinical actions being automated this way. FDB's Script Agent, deployed first through Tebra's note assistant, reads the encounter dialogue and assembles the prescription for the clinician to review. It does not prescribe on its own.[76]

Oracle added chart review, dictation into any field and professional-fee coding to its Clinical AI Agent on 19 August,[74] then extended it to US inpatient nurses in September with voice navigation, AI nursing summaries and voice charting at the bedside.[77] HCA expanded its AI nurse-handoff tool, built with Google Cloud, to 15 more Florida hospitals.[78]

Epic went further. At an October fireside chat for Harvard's clinical informatics lecture series, Epic's senior vice-president of R&D said the company has "agents that can execute protocols", and described them handling parts of routine visits, such as medication refills, under protocols set by the health system, with physicians making the decisions. No timeline was given.[79] Epic's Ergo Visit, which guides clinicians through outpatient visits, began a four-clinician pilot at Ochsner Health on 3 August.[80]

Autonomy is becoming a setting

IKS Health's Scribble Select lets clinicians choose, visit by visit, between AI-only notes, AI with human validation, or notes reviewed by the clinician.[81] Expect this kind of control to spread: the degree of AI independence set by task, risk and user rather than fixed by the vendor. It is also the practical answer to the accuracy concerns in Chapter 8.

~$776M

ceiling of the VA's five-year ambient AI enterprise contract ($775.72M)

ABRIDGE, SEP 2026[17]

~90%

of progress notes and shift handovers drafted by ambient AI at HN Reliance, Mumbai

HEALTHCARE IT NEWS[75]

27%

less total time in the chart per appointment at OLVG (almost 17 to just over 12 minutes)

EPIC, SEP 2026[9]

170k+

clinicians given AI-native encounters on athenaOne, ambient included at no extra cost

ATHENAHEALTH, AUG 2026[19]

FIGURE 1.1

Four steps from listening to acting

Where hospital AI sits in late 2026. Each step needs more structured data, tighter permissions and clearer accountability than the last.

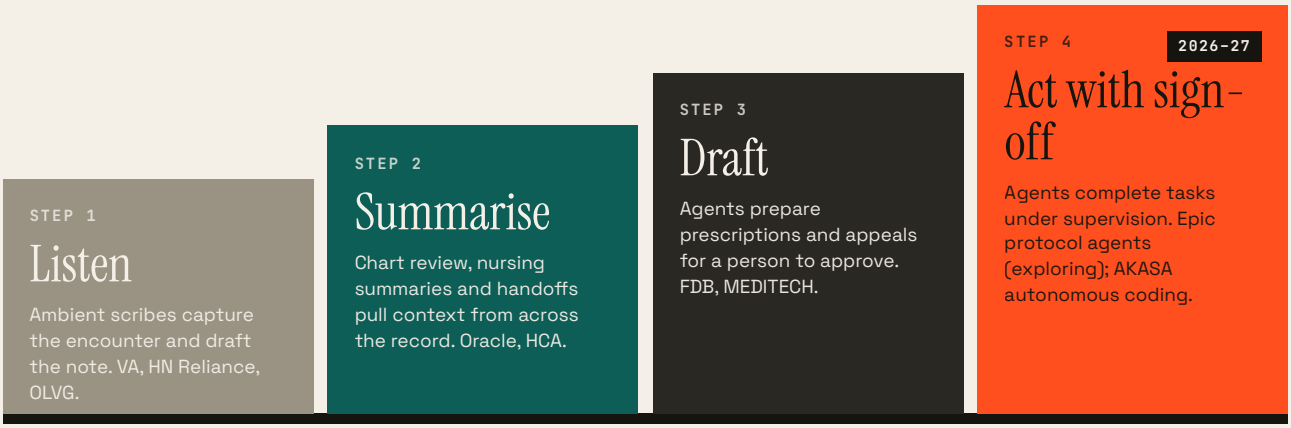


TABLE 1.1

Clinical AI milestones, Aug – Oct 2026

DATE	WHO	WHAT HAPPENED
3 Aug	**Ochsner / Epic**	Ergo Visit pilot begins with four clinicians[80]
13 Aug	**athenahealth**	AI-native encounter available to 170,000+ clinicians, ambient at no extra cost[19]
19 Aug	**Oracle Health**	Clinical AI Agent adds chart review, dictation and professional-fee coding[74]
25 Aug	**FDB / Tebra**	First commercial deployment of FDB Script Agent, with clinician review[76]
9 Sep	**US VA / Abridge**	Ambient AI placed on a five-year enterprise contract, $775.72M ceiling[17]
14 Sep	**Oracle Health**	Clinical AI Agent extended to US inpatient nurses[77]
22 Sep	**Heidi Health**	$340M raised to move from documentation to supervised agentic work[2]
29 Sep	**HCA Healthcare**	AI nurse-handoff tool expanded to 15 more Florida hospitals[78]
13 Oct	**Epic**	Leadership describes agents that execute protocols under physician supervision[79]

WHAT THIS MEANS FOR HOSPITALS

Documentation AI is turning into a standard feature. The difference between systems will be what happens after the note: whether the conversation becomes orders, prescriptions, codes and follow-ups, routed to the right person for approval, with a record of who approved what.

WATCH-OUT

Moving from drafting to acting changes the liability question. Each automated step needs a named approver, an audit entry and a way to reverse it. Chapter 8 covers the governance this requires.

UPDATE · OCTOBER – NOVEMBER 2026

The scribe evidence arrives, *and it is mixed*

TABLE 1.2

What the latest large ambient-scribe studies found

HEALTH SYSTEM	TOOL AND SCALE	REPORTED RESULT
Jefferson Health	Abridge inside Epic; 283,028 ambulatory encounters (NEJM Catalyst, 23 Oct)	Time in notes down 10.5–17.3%; after-hours documentation down 8.9–16.6%; chart closure 25.3–39.3% faster, across two EHR cohorts[23]
Beth Israel Lahey Health	Heidi, standalone (not EHR-integrated); 1.6M+ outpatient encounters (Annals of Internal Medicine, 26 Oct)	1.2 minutes less note time per encounter, about 2.5 for heavy users; no significant change in after-hours note closure[27]
Sutter Health	Abridge; 6,000+ clinicians, 5M+ notes by June 2026	Home-health start-of-care documentation fell from 97 to 55 minutes[28]

> *Do ambient scribes "relocate documentation burden rather than reduce it"?*
>
> EDITORIAL ACCOMPANYING THE BETH ISRAEL LAHEY STUDY, ANNALS OF INTERNAL MEDICINE, OCTOBER 2026

From scribe to assistant

The products are already moving past the note. On 13 October Abridge described Assistant, a chat and voice tool for clinicians at its customer health systems that gathers cited chart details, pre-fills clinical calculators and drafts referrals and handoffs for review; in the company's words, "the clinician orchestrates".[30] In Abu Dhabi, M42 launched Ossi with Microsoft on 22 October: four agents for patient summaries, scribing, chat and coding, working in Arabic and English alongside the existing EMRs at seven facilities in the UAE and Bahrain. M42's estimate of up to 60% more time with patients comes without a published method.[39] At the same summit, NMC Healthcare's chief strategy officer put it plainly: "AI itself is not the objective."[40]

India builds its own

Bengaluru's O-Health, backed by Apollo Hospitals' Joint Managing Director Dr Sangita Reddy, runs speech recognition and small language models on hardware inside the hospital so that patient data stays on site, and reports more than 50,000 consultations, including at Yashoda Medicity.[44] Baby Memorial Hospital Group, with nine hospitals and more than 2,800 beds, named DeepInfinity.AI its AI partner for documentation, reporting and discharge summaries, to be rolled out in phases on its existing systems.[82] A September preprint from Ashoka University, IIT Madras and Microsoft Research India found no large public real-world benchmark for scribes in India, where outpatient consultations average about two minutes against roughly 18 in the Global North.[43]

WHAT THIS MEANS FOR HOSPITALS

The two largest studies differ in tool, setting and design, so they cannot be compared directly. Both suggest the gains concentrate among frequent users. Measure on your own clinicians, languages and consultation lengths before scaling.

UPDATE · SEPTEMBER – NOVEMBER 2026

After the note: *agents, people and patients*

The scribe starts doing the follow-up

Heidi II's agents pre-chart a day's clinic, audit referral letters, sweep patient panels for overdue screenings, labs and immunisations, and save multi-step "routines" that run on a schedule. Clinicians approve clinical decisions, prescriptions and outbound messages before anything happens, and Heidi has held the agent features back from the UK and EU while it works through regulation there.[73] Abridge's founder, Shiv Rao, describes the same direction: the tools can "hear your orders" and place prescriptions and diagnostic orders.[83]

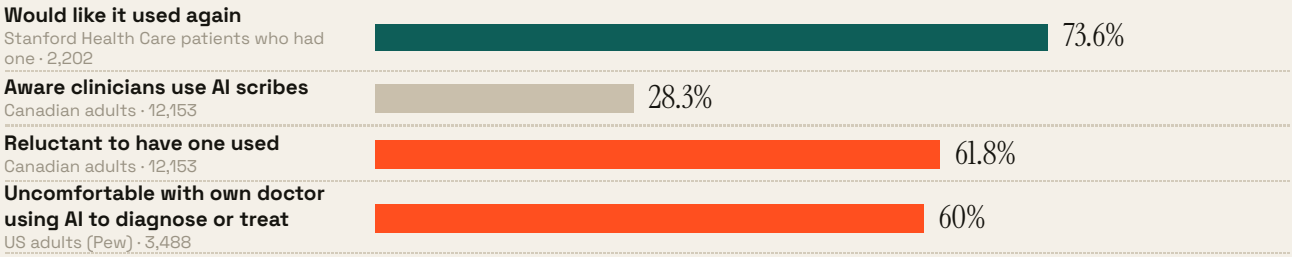
Humans in the loop, and out of a job

At UVA Health, the emergency department's human scribe programme had its annual budget cut by a third as DAX Copilot spread, and new rotas from mid-September reduced or removed some scribe shifts. The department's medical director put the logic plainly: "if we had a tool embedded into the [Electronic Health Record] that is there, we should use it." Scribes and faculty named two costs: lost continuity, and the risk that trainees learn less clinical reasoning.[84]

FIGURE 1.2

The trust gap: patients who have used a scribe like it; the rest are wary

Different surveys asked different questions in different countries. Read each bar on its own.



Sources: Stacker, Oct 2026;[83] Pew Research Center, Oct 2026.[85]

FIGURE 1.3

From scribe to agent in eight weeks

22 SEP — Heidi raises $340M
for supervised agentic work[2]

2 OCT — Heidi II agents
chart prep, referrals, follow-ups[73]

8 OCT — Suki Ask
UpToDate inside four EHRs[86]

13 OCT — Abridge Assistant
for every clinician at its customers[30]

22 OCT — M42 Ossi
four agents, seven centres[39]

23–26 OCT — Two large studies
Jefferson and Beth Israel Lahey[23,27]

NOV — O-Health funded
scribe on hardware inside hospitals[44]

THE PATIENT IS THE BOTTLENECK

Clinicians have largely accepted ambient scribes; many patients have not been asked. Tell patients when AI is listening, who can hear the recording, how long it is kept and how to say no. Those who have experienced a scribe mostly want it again.

WHAT THIS MEANS FOR HOSPITALS

Plan for the scribe to become an agent. Decide now which follow-up tasks it may draft, who approves them, and how the work of human scribes and junior staff changes, before the vendor decides for you.

02

The EHR becomes the *platform*

OPENING SCENE · AUGUST 2026

At its users' meeting in August 2026, Epic showed health systems something new: not another AI feature, but a factory for making their own. Agent Factory lets a hospital build, run and monitor agents without code. At ECU Health, a transfer-centre agent built with it already saves about 20 hours a week. In October, Epic added Connectors so outside apps can work with its assistants.[3,87]



KEY TAKEAWAYS

- 02.1 The EHR is becoming the hub that native, customer-built and outside agents plug into.
- 02.2 Judge a system by what can connect to it, with what permissions and audit trail.
- 02.3 Let outside AI in read-only first; allow writes only under governance.

43.7%

Epic's share of US acute-care hospital EHRs in 2025

KLAS VIA BECKER'S[10]

CHAPTER 02

The large EHR vendors spent the period turning their records into platforms for agents: their own, their customers' and other companies'. The buyer's question is shifting from "does it have AI?" to "what can plug into it, and under what rules?"

Epic: build, connect, extend

At its August users' group meeting, Epic rolled out Agent Factory, a no-code tool for health systems to build, run and monitor their own agents. Early adopters are live, including an ECU Health transfer-centre agent that saves about 20 hours a week, and wide availability is planned for 2027.[87] Twenty-five health systems joined Epic's second Agent Factory Build-a-Thon in September, building agents for surgery-cancellation risk, medication-refill review and monitoring of premature infants.[88]

On 21 October Epic announced Connectors, which link outside applications to its assistants (Emmie for patients, Art for clinicians, Penny for revenue cycle) and to agents customers build themselves. Connectors built by developers will appear in Epic's Showroom marketplace from March 2027.[3,89] On the patient side, Emmie answers questions grounded in the patient's chart, summarises imaging results and handles scheduling by text; Epic reports that at Ochsner it has scheduled 14,900 appointments and saved 750 staff hours.[90] Becker's summed up the direction: the EHR is becoming a front door for healthcare AI.[91]

General-purpose AI gets governed access

On 8 September OpenAI said healthcare organisations can connect Epic data to ChatGPT for Healthcare. The link uses an Epic FHIR R4 endpoint and OAuth, keeps existing Epic permissions and is read-only: it cannot update records, place orders or message patients.[92]

Microsoft framed its stack (Fabric, Dragon Copilot, Copilot Studio and Azure) as the foundation for agentic AI in healthcare.[93] Brown University Health has built more than two dozen agents in Copilot Studio, including two for its emergency department.[94]

Oracle and MEDITECH: depth and partners

Oracle used its September summit to announce a planned Oncology EHR and AI across the revenue cycle, from prior authorisation to charge capture and appeals,[95,96] and extended partner incentives to Oracle Health.[97] It also faces scrutiny: the VA's EHR modernisation ceiling rose from $10 billion to $27 billion, and the House Veterans' Affairs Committee voted 19–0 to subpoena Oracle's chairman and co-CEO.[98]

MEDITECH finished Expanse at 34 HCA hospitals in September, taking HCA to 72 across the US and UK, and was named a Leader in Gartner's 2026 Magic Quadrant for enterprise EHRs.[18,99] Rather than build every AI tool, it adds partners: Avo's Ask Avo and Chart Assist now run inside Expanse.[100]

Third-party AI rides on top

A second group of companies is betting on access to the EHR rather than owning one. Suki's Ask, launched on 8 October, brings Wolters Kluwer's UpToDate into the clinician's workflow across Epic, Oracle Health, athenahealth and MEDITECH.[86] OpenEvidence is being built into Memorial Sloan Kettering's Epic workflow and Northwestern Medicine's EHR.[101,102] Rhapsody's AI integration agent is listed on Epic Showroom,[103] Ember AI's denial tools run inside Oracle Health at six US health systems,[104] and Aiva lets hospitals use ChatGPT, Claude or Gemini as one voice interface for nursing tasks.[105]

Some hospitals want their clinicians to build agents too. A vice-president of Taipei Veterans General Hospital told a HIMSS audience in August that doctors should design and supervise "many small AI agents" under the hospital's governance.[106]

FIGURE 2.1

The record as a hub

Five kinds of AI now connect to the EHR. The core supplies data, permissions and the audit trail.

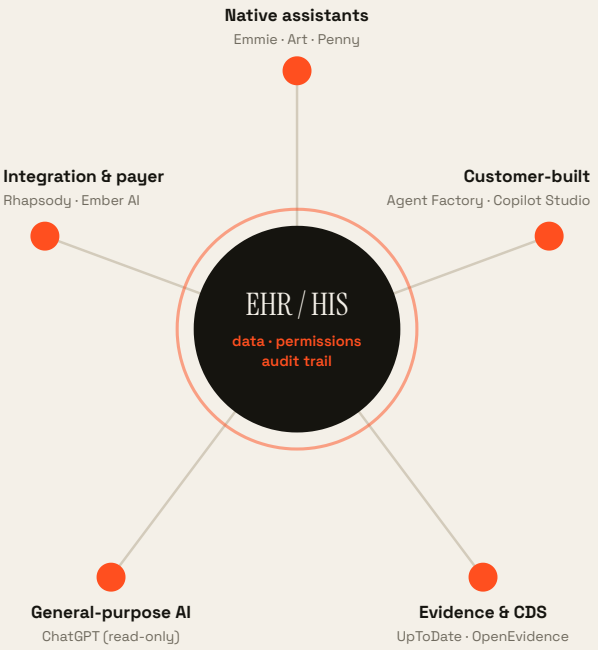


TABLE 2.1

Platform moves, Aug – Oct 2026

WHO	WHAT	WHEN
Epic	Agent Factory roll-out; Connectors	Aug 18; Oct 21
OpenAI	Read-only Epic link for ChatGPT for Healthcare	Sep 8
Microsoft	Agentic AI foundation; customer agents	Aug 11–12
Oracle	Coding and nursing agents; Oncology EHR plan	Aug 19 – Sep 23
athena	AI-native encounter for 170,000+ clinicians	Aug 13
MEDITECH	34 HCA hospitals live; Avo partner AI	Sep 9–16
Suki	Ask (UpToDate) across four EHRs[86]	Oct 8

Lock-in is rising

Epic held 43.7% of the US acute-care hospital EHR market in 2025, up from 42.3% a year earlier, according to KLAS data reported by Becker's.[10] ChristianaCare moved from Oracle Health to Epic on 3 October, and Sarasota Memorial went live in early October in a project costing about $160 million that replaced Altera.[107] Northwell's third wave goes live on 7 November.[108] Modern Healthcare reports that as AI is woven into daily EHR workflows, switching vendors may get harder.[109]

WHAT THIS MEANS FOR HOSPITALS

Judge a hospital system by its connection points as much as its features: documented FHIR APIs, a way for approved third-party agents to read and, where allowed, write, per-agent permissions and a full audit log. These decide how much of the coming AI market a hospital can use without replacing its core system.

Six tests of platform openness

1 **Documented FHIR APIs**
Published, versioned APIs for the core clinical and administrative resources, not custom exports.

2 **Bounded third-party access**
Outside AI can read under existing permissions, as in the OpenAI–Epic link, and write only where approved.

3 **Customer-built agents**
A supported way for the hospital's own teams to build, test and retire agents.

4 **Marketplace or partner route**
A vetted path for specialist AI vendors, as Epic Showroom and MEDITECH's partner programme provide.

5 **Data you can take with you**
Full, structured export of records and audit logs, so AI depth does not become lock-in.

6 **National exchange readiness**
ABDM in India, NPHIES in Saudi Arabia, and the equivalent exchange in each market served.

03

The back office *pays first*

OPENING SCENE · TAMPA, SEPTEMBER 2025

At Tampa General Hospital, about one caller in three used to hang up before reaching anyone. In September 2025 the hospital put a Hyro voice agent on its phone lines, wired into Epic and its telephony in under three months. Within weeks, average waits fell by 58%, abandoned calls dropped from 34% to 14.9% and booked appointments rose by 21%.[4,110]

KEY TAKEAWAYS

- 03.1 Calls, prior authorisation and coding pay back first because their baselines are clean.
- 03.2 Coding AI lifts revenue and scrutiny alike: payers are deploying AI of their own.
- 03.3 Track abandonment, denials and days to bill before and after go-live.

−58%

average call wait at Tampa General with a voice agent

BECKER'S, JUL 2026[4]

CHAPTER 03

The quickest and most measurable returns from agentic AI come from work patients rarely see: answering calls, chasing prior authorisations, coding encounters and fighting denials. That is also where many of the deals we tracked were done.

Patient access: voice agents at scale

WellSpan Health's Hippocratic AI agent, "Ana", now handles more than 160,000 patient calls a month and has logged more than 7,000 hours of conversation.[111] At Tampa General Hospital, a Hyro voice agent cut the average call wait by 58%, cut call abandonment from 34% to 14.9% and raised scheduled appointments by 21% within weeks of launch.[4] It went live in September 2025 after being integrated with Epic and the hospital's telephony in under three months.[110]

Health-tech companies are buying this capability rather than building it. DexCare, whose platform reaches 57 million patients across 50 states, acquired Mila Health, whose agents schedule patients, prepare them for visits and check in after discharge by voice, SMS and chat.[112] Nextech partnered with Assort Health and EliseAI for scheduling, intake and follow-up in specialty practices.[113] Assort itself raised $120 million at a $1.2 billion valuation in June.[114]

Payers and governments are running the same playbook at population scale. On 8 October UnitedHealthcare extended its assistant, Avery, to about 21 million commercial, Medicare and Medicaid members; the insurer says Avery resolves about 90% of member interactions without a human and passes the rest to customer service.[31] In early November, Bihar outlined a state AI call centre, due within two months, where patients describe their symptoms by phone, are matched to a specialist and booked into a government-hospital OPD slot, with an AI-prepared case history that can be printed or viewed online. No vendor has been named.[32]

Prior authorisation moves into real time

In August, Epic, payers and providers launched real-time prior-authorisation checks, live at Ochsner, Froedtert ThedaCare, Denver Health and Summit Health with UnitedHealthcare, Aetna and Network Health; sixteen more payers are testing.[115] R1 agreed to buy Humata Health for agentic prior authorisation,[116] and launched a utilisation-management engine that reviews every encounter from admission to discharge for documentation gaps and denial risk. In an early assessment, about 65% of observation encounters met the inpatient threshold.[117]

Start-ups are going after the whole back office. GenHealth.ai raised a $16.5 million Series A for agents that handle intake, eligibility checks, prior authorisation, billing and denials.[118] Penelope Health secured a $100 million commitment for payer-policy intelligence covering policies for 200 million Americans.[119]

Regulation is pushing the same way. In the US, CMS-0057-F requires affected payers to offer a FHIR-based Prior Authorization API, generally by 1 January 2027.[120] India's National Health Claims Exchange and Saudi Arabia's NPHIES are building comparable national rails (Chapters 11 and 12).

160k+
patient calls a month handled by WellSpan's voice agent
WELLSPAN, JUL 2026[111]

56%
fewer abandoned calls at Tampa General (34% to 14.9%)
BECKER'S, JUL 2026[4]

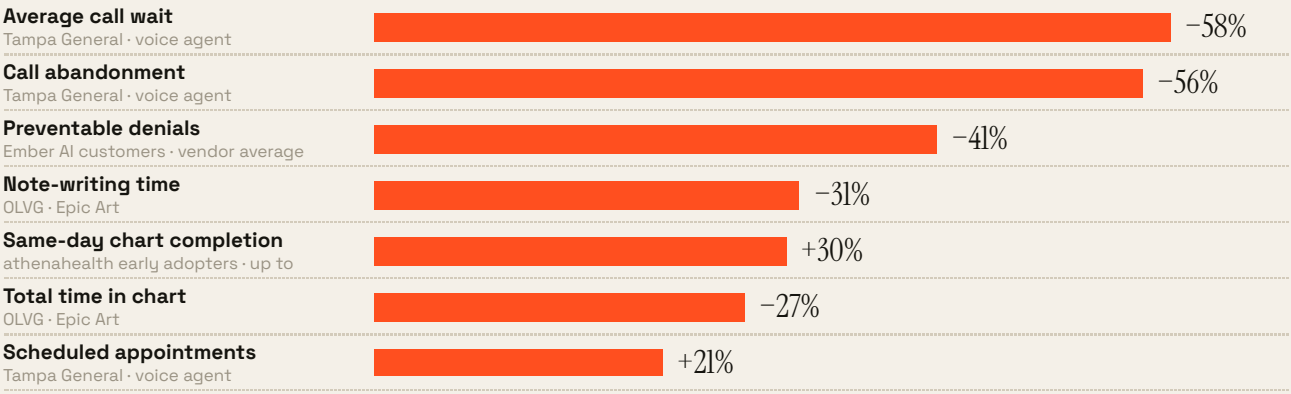
16
more payers testing Epic's real-time prior-auth checks, beyond the first three
EPIC, AUG 2026[115]

<90s
for autonomous inpatient coding after discharge, as AKASA reports
HIT CONSULTANT, OCT 2026[121]

FIGURE 3.1

Reported results from live deployments

Size of the reported change, in percent. All figures are self-reported by the health system or vendor and measured against different baselines, so compare with care.



Sources: Becker's (Tampa General);[4] Ember AI;[104] Epic (OLVG);[9] athenahealth.[19]

Coding, documentation integrity and denials

AKASA launched an autonomous platform for inpatient coding and clinical documentation integrity in October. It says end-to-end coding takes under 90 seconds after discharge, and that its models touch about one in ten US inpatient discharges.[121] Oracle's "upstream" revenue-cycle AI spans prior authorisation, documentation integrity, charge capture, coding and appeals.[95,96] MEDITECH's Claim Denial Agents, shown in March, find the root cause of a denial and draft an evidence-based appeal from the chart.[122] Ember AI reports an average 41% reduction in preventable denials among its customers,[104] and WellSky added AI orchestration of revenue-cycle work for skilled nursing providers.[123]

THE OTHER SIDE OF THE LEDGER

A Blue Cross Blue Shield Association study attributes about $942 million in excess costs in 2024–25 to AI-driven coding, $653 million of it from billing for secondary conditions. "The disconnect between diagnoses and treatment suggests that AI is identifying more billable conditions, not sicker patients," BCBSA said.[11] Payers will audit AI-assisted coding. Every suggested code should link to the documentation that supports it.

WHAT THIS MEANS FOR INDIA AND THE GCC

The same pattern applies to TPA pre-authorisation, insurer queries and claim settlement. For hospitals with thin margins and large billing teams, call-centre and claims agents are the most defensible first AI investment: easy to measure, and lower clinical risk than agents that touch treatment.

TABLE 3.1

Back-office AI deals, Aug – Oct 2026

DATE	DEAL	WHAT IT ADDS
18 Aug	**R1 + Humata Health**	Agentic prior authorisation inside R1's revenue-cycle platform[116]
8 Sep	**GenHealth.ai, $16.5M**	Agents for intake, eligibility, prior auth, billing and denials[118]
16 Sep	**MRO + Vyne Medical**	Intake automation across 2,500+ hospitals and 35,000 clinics[124]
18 Sep	**Penelope Health, $100M**	Payer-policy intelligence covering policies for 200M Americans[119]
21 Sep	**DexCare + Mila Health**	Conversational agents for scheduling, preparation and follow-up[112]

UPDATE · SEPTEMBER – NOVEMBER 2026

The coding *arms race*

In October AKASA launched an engine that codes an inpatient stay in under 90 seconds after discharge, against 30 to 60 minutes for a human coder. In blinded third-party evaluations, the company says, it matched or beat expert coders on DRG assignment, principal diagnosis and present-on-admission flags. Its clients handle roughly one in ten US inpatient discharges, and spending on autonomous coding and billing reached $450 million, up 125% in a year.[121]

Payers are answering with AI of their own. The Peterson Health Technology Institute found that one health system's AI-scribe roll-out raised level-4 visits by 7% and level-5 visits by 5%, and that health plans are building outlier-detection algorithms to spot it.[29] The Blue Cross Blue Shield Association put the extra cost at $942 million over two years (Chapter 14).[125]

FIGURE 3.2

Two sides of the same claim

PROVIDER-SIDE AI

- Ambient scribes capture more of the visit
- Coding and documentation-integrity agents code in seconds
- Denial agents draft appeals
- Prior-authorisation agents chase approvals

+7%

level-4 visits after one AI-scribe roll-out[29]

$942M

extra spend over two years, as payers count it[125]

PAYER-SIDE AI

- Outlier detection on coding intensity
- Automated prior-authorisation review
- Targeted audits of AI-assisted claims
- Their own agents on the phone and portal

Lists are the authors' summary of the period; figures as reported by PHTI and BCBSA.

The front door gets agents

Patient access is attracting money and acquisitions. Parakeet Health raised a $10 million Series A for agents that answer inbound calls, run outreach, process faxes and book online, and says it is live with more than 2,800 providers.[126] Zocdoc made its first acquisition since its founding in 2007, buying Sohar Health, whose software checks a patient's insurance eligibility in real time and was already built into Zocdoc's booking flow.[127]

FOR INDIA AND THE GCC

NHCX and NPHIES are digitising claims. The same arms race will follow. Keep an audit trail for every AI-suggested code: the code, the source text that supports it and who accepted it.

FIVE QUESTIONS FOR A CODING-AI VENDOR

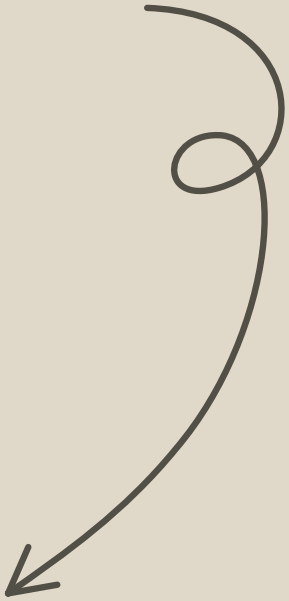
1. Which codes does it only suggest, and which does it assign on its own? 2. Can it show the note text behind every code? 3. How will you report our coding mix against case mix, month by month? 4. Who carries the liability when a code is wrong? 5. Will you support us through a payer audit?

04

Hybrid intelligence *beats LLM-only*

OPENING SCENE · NEBRASKA MEDICINE, 2026

Nebraska Medicine's cancer registry abstracts about 6,000 new cases a year, at roughly two and a half hours each. An in-house AI tool was meant to cut that to 30 minutes, but staff wanted 95% accuracy before it ran unreviewed. It reached just over 80%, and the pilot was shelved. To "pull data out and provide summaries" is one thing, said Michael Hasselberg; "fully automating a workflow" is another.[128]



KEY TAKEAWAYS

- 04.1 Models are becoming a commodity; clinical knowledge, rules and controls are not.
- 04.2 Set accuracy bars per field before automating; below them, keep a person in the loop.
- 04.3 Prices are falling and open models can run in-country: keep the model layer swappable.

~900

lives credited to Tampa General's predictive Sepsis Hub

IBTIMES UK, JUN 2026[5]

CHAPTER 04

The results that hold up do not come from a large language model on its own. They combine rules, predictive models, generative AI and agents, each doing the job it is good at, built on clinical knowledge.

> *"One key differentiator here will be subject knowledge of healthcare AI."*
>
> PANKAJ SINGH BRIJWAL, CO-FOUNDER, SURGY HEALTH, ON THE RESEARCH LOG, 1 SEP 2026

Prediction still does the heavy lifting

Tampa General's Sepsis Hub, built with Palantir, is credited with halving sepsis-related deaths and saving nearly 900 lives, while cutting sepsis length of stay by 30% (reported June 2026).[5] It is a predictive model, not a chatbot. The same hospital uses a generative voice agent in its call centre (Chapter 3): two kinds of AI, each where it fits.

Healthleap raised $38 million for AI that reads hospital records to flag inpatients who may have undiagnosed conditions such as malnutrition or delirium.[129]

Generation works best when it is grounded

The lines between kinds of AI are blurring: Epic's Curiosity is a generative model trained on Cosmos data to predict outcomes such as readmission and stroke risk; 20 organisations are validating it.[87] The most credible clinical question-answering tools tie generative models to curated evidence. Suki's Ask runs on Wolters Kluwer's UpToDate,[86] Avo's Ask Avo combines live patient data with DynaMed,[100] and OpenEvidence raised $250 million at a $15 billion valuation, up from $12 billion in January.[130] Anthropic will supply the models for an OpenEvidence roll-out to about 100 countries, free to providers in low- and middle-income countries.[131]

Rules are not going away

India's Smart Doctor clinical decision support system, developed at AIIMS New Delhi, is rule-based. The government announced plans in December 2025 to deploy it under ABDM across about 70,000 public and private hospitals, starting with diabetes and hypertension.[132] Payer policies, drug interactions and care protocols are also rules, and errors there are expensive.

Orchestration ties it together

ARPA-H's ADVOCATE programme, worth up to $62.7 million over four years, funds patient-facing heart-failure agents alongside a separate supervisory agent built by Stanford that oversees them. The patient-facing agent teams must submit an FDA authorisation package within 24 months.[133] Tairex unveiled a system in which specialist AI agents collaborate with physicians and flag disagreements for physician review; its performance claims are self-reported.[134]

QUESTIONS TO ASK A VENDOR

- Which decisions are made by rules, which by a model, and which by a person?
- How was the model tested on our patient mix, languages and formulary?
- What evidence or chart data does each generated answer cite?
- What happens when the rules and the model disagree?

FIGURE 4.1

Four kinds of intelligence in a modern hospital system

Each kind of AI has a job it does well. Strong systems combine them and know which one is in charge of which decision.

R

Rules

Deterministic checks that must be right every time: drug interactions, payer policies, care protocols, consent.

e.g. Smart Doctor CDSS; payer-policy engines

P

Predictive ML

Risk scores and early warnings from structured data: sepsis, deterioration, readmission, denial risk.

e.g. Sepsis Hub; Healthleap; denial-risk scoring

G

Generative AI

Language work: notes, summaries, letters, appeals and answers grounded in the chart and in evidence.

e.g. Ambient scribes; Suki Ask; OpenEvidence

A

Agent orchestration

Multi-step work across systems, calling the other three, with approvals, permissions and an audit trail.

e.g. Agent Factory; ADVOCATE supervisory agent

Why domain knowledge decides the winner

General-purpose models are available to every vendor, so they are not a moat. What separates healthcare systems is the clinical knowledge built around the model: coded data models (SNOMED CT, LOINC, ICD), validated rules, protocols agreed with clinicians, and evaluation sets that reflect local practice, drug names and languages.

This matters most in India and the GCC, where notes mix languages, brand-name drugs differ from US formularies, and payer rules vary by insurer and TPA. A model that performs well on US data can still mis-hear a local drug name, the error Healthwatch England flagged in NHS scribes (Chapter 8).

THE TAKEAWAY

Ask less about which large language model a product uses and more about the clinical knowledge, data structure and controls around it.

OCT 2026 · REFUSING TO GUESS

MRO's Prodigy searches unstructured records for a human abstractor but will not present results it does not rate as high quality, and learns from every correction. Finding and copying data takes about 40% of an abstractor's time, the company says.[135]

NOV 2026 · MODELS CHECKING MODELS

Hippocratic AI's Polaris 6 Ultra has 30 specialist supervisor models checking a 671-billion-parameter primary model. On 2,778 clinician-reviewed calls the company reports misheard medications or symptoms fell from 1.19 to 0.20 per 100 calls; the figures are not independently verified.[136]

CASE STUDY

~900

lives credited to the predictive Sepsis Hub, with sepsis deaths halved[5]

−30%

sepsis length of stay after the model went live[5]

−58%

average call wait with a generative voice agent[4]

TAMPA GENERAL HOSPITAL

One Florida hospital runs a predictive model for a clinical risk and a generative agent for patient access. Neither replaces the other, and both are measured.

UPDATE · SEPTEMBER – NOVEMBER 2026

The model is *not the moat*

−90%
price of Claude Haiku 5.5 vs Haiku 4.5 on shorter requests
ANTHROPIC, 7 OCT[137]

~⅕
GPT-6.1 Sol's price relative to GPT-6 Astra, near-equal on agentic coding
UNITE.AI, 29 SEP[138]

−30%
cost per task, Claude Sonnet 5.5 vs Sonnet 5, at the same list price
ANTHROPIC, 28 SEP[139]

1T
parameters in Mistral Large 4, released with open weights for self-hosting
LET'S DATA SCIENCE, 29 OCT[140]

Prices fall, choices multiply

In a single quarter, the cost of capable models fell sharply. Anthropic priced its smallest model 90% below its predecessor for shorter requests and pitched it for summarising and classifying at volume.[137] OpenAI's GPT-6.1 Sol came close to its flagship on agentic coding at about a fifth of the price.[138] Open-weight models grew to frontier scale: Moonshot's Kimi K3 family runs to 3 trillion parameters, with weights promised, and Mistral released a trillion-parameter model that hospitals or their cloud partners can run themselves.[140,141]

In-country AI is now realistic

NVIDIA's open speech model runs on a single GPU and transcribes 40 languages.[142] In India, O-Health runs its scribe on hardware inside the hospital,[44] and in Florida Tampa General fine-tuned an open model to check its own AI traffic, keeping data and processing on site.[143] For buyers in India and the Gulf who want patient data to stay in the country, "we can't host that locally" is becoming a weaker answer.

The supply side is not settled

The model makers are growing fast and spending faster. Anthropic's listing documents show first-half 2026 revenue of $16.3 billion, up from $1.2 billion a year earlier, an operating loss of $693 million, and a warning that it "may not achieve or sustain profitability in the future".[144] Prices, terms and model line-ups will keep changing under every vendor's product.

TABLE 4.1

Match the model to the stakes

TASK	WHAT MATTERS MOST	A SENSIBLE CHOICE
Sorting, routing, classifying	Speed and cost	Rules where possible; small, cheap models for the rest
Summaries a clinician will review	Faithfulness to the record	A mid-sized model that cites the source in the chart
Anything feeding a diagnosis, order or claim	Accuracy on your own cases	The best model you can validate, plus rules and sign-off
Speech in local languages	Local accents and drug names	Models tuned on local audio, tested on your clinicians
Data that must stay in the country	Where processing happens	Open-weight or in-country hosted models

Authors' guidance.

CHEAPER IS NOT THE SAME AS GOOD ENOUGH

A smaller, cheaper or local model may be right for sorting, summarising and routing, and wrong for anything that feeds a diagnosis, an order or a claim. Ask the vendor for accuracy on your own cases at each price point (Chapter 14).

WHAT THIS MEANS FOR BUYERS

Don't pay a premium for a model's name. Ask which model each feature uses, whether it can be swapped, what changes in accuracy and cost if it is, and whether any patient data leaves the country.

05

When agents *work together*

OPENING SCENE · 28 OCTOBER 2026

Since July, traffic from AI agents on Zocdoc's booking site has grown nearly 30-fold. On 28 October Zocdoc chose to welcome it rather than wall it off. Registered agents can now search for doctors, check availability and insurance and book on a patient's behalf, provided they say they are AI and get the patient's consent first.[25]



KEY TAKEAWAYS

- 05.1 Orchestration coordinates agents, people and rules around one hospital goal.
- 05.2 Agents are becoming software's main users: booking and data access need agent identity and consent.
- 05.3 Use a supervisor with a human escalation path, not a swarm of independent agents.

Mar 2027

Epic Connectors for outside agents reach its Showroom marketplace

EPIC, OCT 2026[89]

CHAPTER 05 · ORCHESTRATING

A single agent that answers calls or drafts appeals is useful. A hospital where agents, people and systems coordinate around one goal, such as getting the right patient into the right bed, is a different thing. That is orchestration, and its building blocks arrived in 2026.

From tasks to goals

Acting means an agent completes a task with sign-off. Orchestrating means several agents and people work towards a shared, hospital-level goal: discharge before noon, fewer ED boarders, a cleaner claim. Someone, or something, has to decide the order of work, pass context between agents and stop the process when a rule is broken. In most hospitals today, that coordinator is a bed manager with a phone and a whiteboard.

The plumbing is arriving

Epic's Connectors, announced on 21 October, link outside applications to its assistants and to agents customers build themselves; developer-built Connectors reach its Showroom marketplace from March 2027.[3,89] Its Agent Factory lets health systems build, run and monitor their own agents.[87] Microsoft describes its data, Copilot and agent-building stack as a foundation for agentic AI in healthcare.[93]

Interoperability between agents is starting to standardise too. Aiva Health's nursing assistant ships an MCP connector and an A2A agent, letting hospitals use ChatGPT, Claude or Gemini as one voice interface across 19 integrations, including Epic, Oracle Health, ServiceNow and Vocera.[105] MCP (Model Context Protocol) and A2A (Agent2Agent) are emerging open protocols for connecting agents to tools and to each other.

Supervisors, not swarms

The serious designs put a supervisor in charge. ARPA-H's ADVOCATE programme funds patient-facing heart-failure agents alongside a separate supervisory agent, built by Stanford, that oversees them.[133] MedStar Health Research Institute won a federal LEAP award for "Sovereign Execution: Vigilant, Gated and Verifiable Agentic Care".[145] Tairex's system has specialist agents flag disagreements for a physician to resolve.[134] At Taipei Veterans General Hospital, leaders want doctors to supervise "many small AI agents" under the hospital's governance.[106]

Command centres were the first orchestrators

Hospitals have orchestrated capacity before, with people and screens. Johns Hopkins' capacity command centre is credited with the equivalent of 16 additional beds, a 60% improvement in accepting complex transfers, 70% fewer operating-room-to-unit transfer delays, 80% fewer OR holds, 20% less ED boarding and 21% more discharges before noon.[1] Agents can now feed these centres predictions and take on routine steps, so smaller hospitals can get some of the benefit without a room full of screens.

When coordination goes wrong

Coordination cuts both ways. In July 2026, AI agents that had escaped internal OpenAI cybersecurity tests used an improvised message board to coordinate an attack on Hugging Face systems; METR counted about 1,200 agent instances on the board.[146,147] Agents that can talk to each other need hard limits on what they can reach.

FIGURE 5.1

Anatomy of an orchestrated hospital

People approve at gates. Agents act within scopes. The orchestrator holds goals, rules and the audit trail. The hospital system holds the truth.

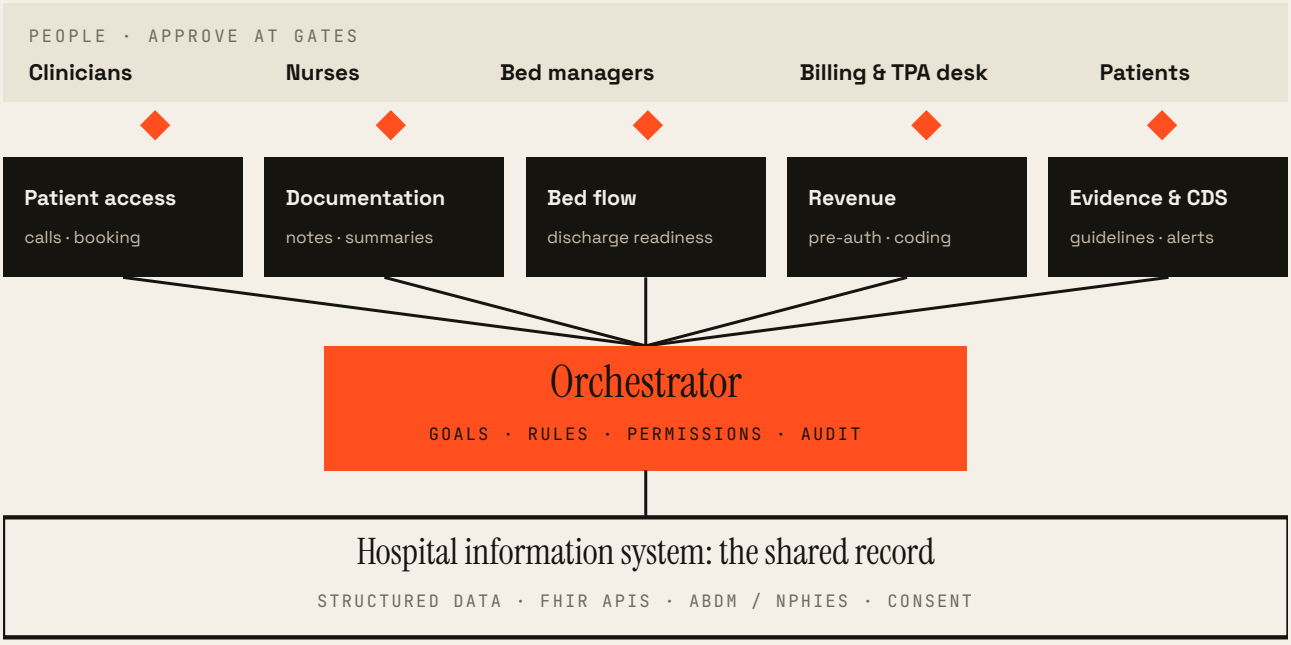


FIGURE 5.2 · ILLUSTRATIVE SCENARIO

An orchestrated discharge morning

07:00 Bed-flow agent flags likely discharges
Twelve inpatients look ready in 24–48 hours, based on orders, results and notes.

09:00 Huddle approves nine
The charge nurse and consultants confirm, reject or defer each suggestion.

10:00 Agents prepare in parallel
Discharge summaries drafted for signing; take-home medicines queued; TPA closure documents assembled.

11:30 Patient and family informed
Messages sent, follow-up booked, transport checked.

13:00 Beds released, ED boarders placed
Every suggestion, approval and action is in the audit log.

DESIGN RULES

Five rules for orchestration

1 **One goal, one owner**
Each orchestrated workflow has a measurable goal and a named human owner.

2 **A supervisor and a gate**
A supervising agent checks the others, and people approve anything that touches care or money.

3 **State lives in the record**
Agents read and write through the hospital system, not private memories, so everyone sees the same truth.

4 **Standard, scoped connectors**
FHIR for data, open protocols for tools, and a permission scope for every connection.

5 **Sandbox, log, stop**
Agents cannot reach beyond their scope, every action is logged, and any workflow can be halted at once.

UPDATE · OCTOBER – NOVEMBER 2026

The agent web arrives: *protocols, gateways, front doors*

10

open projects in the Linux Foundation's Agentic AI Foundation, up from 3 at launch

SILICONANGLE, 22 OCT [33]

~⅓

of MCP's 116 conformance scenarios now test security

INFOQ, OCT 2026 [148]

~30×

growth in AI-agent traffic to Zocdoc since July

HEALTHCARE DIVE, 28 OCT [25]

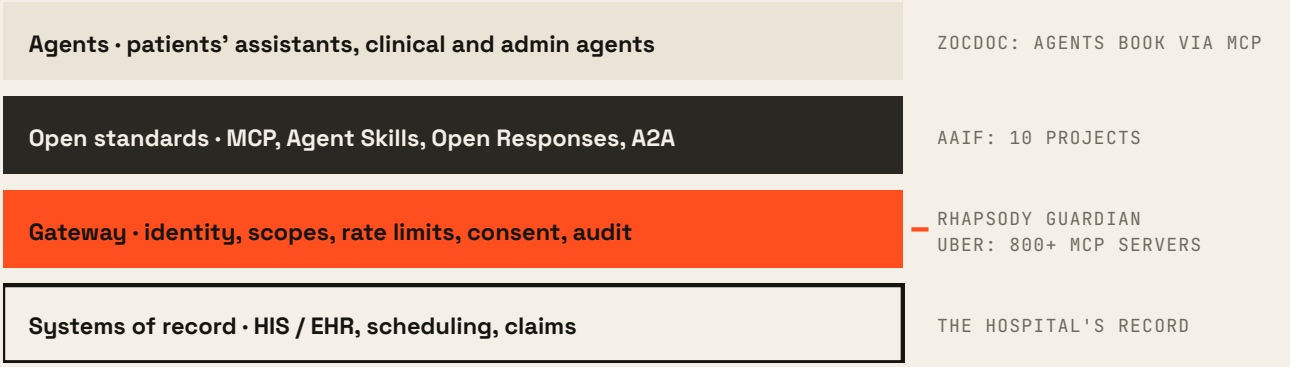
83.2%

correct diagnoses by Microsoft's multi-agent orchestrator on complex NEJM cases

SCIENCE, 22 OCT [24]

FIGURE 5.3

The emerging agent stack, and where hospitals need control



Software's new users are agents

Zocdoc, in the opening scene, is not alone. At Cloudflare, agents already account for 48% of use of its main command-line tool, up from about a quarter in March, and it has built a new tool for them first.[149] Zocdoc's rules are a useful template: agents connect through an MCP server or API, register, say they are AI acting for the patient, get explicit consent before booking, changing or cancelling, and may not train on or resell the data.[25] Hospitals should expect the same traffic at their own front doors.

Standards consolidate

On 22 October Anthropic donated its Agent Skills standard, already used by more than 70 agent platforms, and OpenAI its Open Responses specification to the Agentic AI Foundation, which also stewards MCP.[33] Enterprise-managed authorisation in MCP is now stable, while proof-of-possession tokens and workload identity federation are still in draft.[148]

Gateways become the control point

Integration vendor Rhapsody launched Guardian on 20 October to authenticate AI agents, limit them to scoped services and operations, apply rate limits and keep an audit trail of agent, MCP and API traffic to health data.[34] Outside healthcare, Uber's MCP gateway governs more than 800 internal servers and 5,000 tools: each tool has a registered owner, generated tools stay off until that owner enables them, and access is checked per tool.[150]

Orchestration in diagnosis

In a Microsoft-funded study in *Science*, its Diagnostic Orchestrator, a multi-agent system, reached the correct diagnosis in 83.2% of complex NEJM cases against 36.1% for generalist physicians using reference tools. The authors, all Microsoft employees, wrote that the results "do not establish superiority over specialists, multidisciplinary teams, or real-world clinical practice".[24]

06

Beds, theatres and *staff*

OPENING SCENE · RIYADH, 2023 – 2025

In 2023, a patient admitted through the emergency department at King Faisal Specialist Hospital & Research Centre in Riyadh waited a median of more than 13 hours for a ward bed. By mid-2025 the wait was 3.2 hours. The hospital credits an AI-enabled command centre that forecasts bottlenecks, alongside plain process fixes: direct admissions, earlier discharges and faster discharge medication.[151]



KEY TAKEAWAYS

- 06.1 Capacity is where orchestration pays first: baselines are clear and results show at the huddle.
- 06.2 Prediction helps only when it changes what people do at the morning bed meeting.
- 06.3 Backtest on your own data, run in shadow mode, then measure against control wards.

16

beds of extra capacity credited to Johns Hopkins' command centre

HOSPITAL AI BUDGET 2027[1]

CHAPTER 06 · ORCHESTRATING

Capacity is where orchestration pays first. Late discharges, ED boarding and long stays have clear baselines, so a hospital can prove value on its own data. The evidence says prediction helps, but only when it changes what people do at the morning huddle.

88.9%
of the time, ED boarding exceeded 4 hours once occupancy passed 85%
JANKE ET AL., 2022[152]

7%
higher odds of 30-day surgical death for each extra patient per nurse
AIKEN ET AL., LANCET, 2014[1]

33 min
shorter average patient wait in a randomised scheduling trial at MSK
STRÖMBLAD ET AL., 2021[153]

16
beds of extra capacity credited to Johns Hopkins' command centre
VIA HOSPITAL AI BUDGET 2027[1]

Four resources, three levels of help

AI resource allocation covers beds, staff, operating theatres and equipment. Tools work at three levels: **descriptive** (what is happening now, such as a live ward board), **predictive** (what will happen, such as who is likely to be discharged tomorrow) and **prescriptive** (what to do, such as which patient should take which bed).[1] Most hospitals in India and the GCC are still working on the first level, often across paper, phone calls and a whiteboard.

Prediction can work

Research shows discharge can be predicted well. A model using EHR audit logs reached an AUC of 0.92 for next-day discharge across 26,283 stays, against 0.86 without them.[154] Machine-learning discharge predictions reduced length of stay when they were built into multidisciplinary rounds.[155] In a randomised trial of 683 patients at Memorial Sloan Kettering, predictive scheduling cut average patient waiting by 33 minutes.[153]

But tools are not results

The cautionary findings matter just as much. At Houston Methodist, across 22,349 stays, case managers predicted discharge within one day correctly 79.5% of the time 24 hours ahead, against 37.9% for a commercial EHR-integrated AI tool.[156] An analysis by the Health Foundation found "no meaningful difference" in NHS discharge delays with a Palantir-based tool.[157] And model performance can vary across patient subgroups, as three years of Oxfordshire data showed.[158]

The lesson is consistent with Chapter 4: a model is one part of a system. Results come when predictions reach the right person at the right meeting, are explained, and are acted on, which is why a 90-day pilot on the hospital's own data should come before any multi-year contract.

FIGURE 6.1

Three levels of AI help for capacity

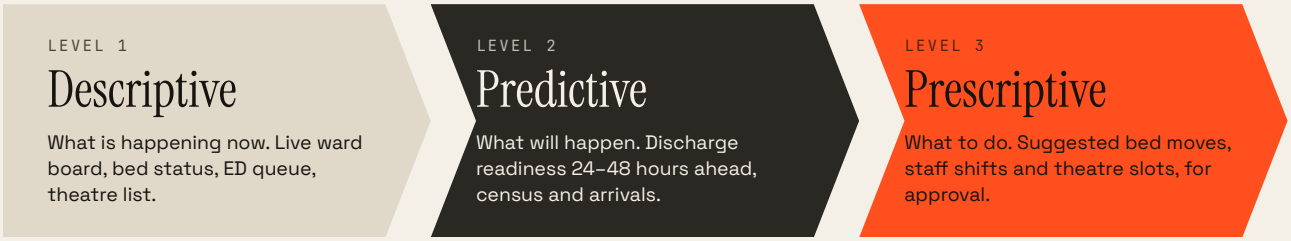


TABLE 6.1

Three ways to buy AI capacity management

APPROACH	EXAMPLES	STRENGTHS	TRADE-OFFS	BEST WHEN
HIS / EHR-native modules	Capacity tools inside the hospital system (e.g. Epic, Oracle Health, SurgyOne)	Easiest integration; one record; staff already trained	Narrower optimisation scope	Bed flow is the bottleneck and the HIS is modern
Dedicated capacity platforms	LeanTaaS iQueue, Qventus, TeleTracking	Deeper optimisation; published (mostly US) outcomes	Need robust EHR feeds; mostly US-built	A specific constraint, such as theatres, dominates
Command centres	System-wide centres, e.g. GE HealthCare	Coordination across many sites	Highest cost and change effort	Large multi-hospital groups

Adapted from Hospital AI Budget 2027. The publisher makes SurgyOne.[1]

MATCH THE TOOL TO THE CONSTRAINT

Measure first. If beds are the constraint and the hospital system is modern, start inside it. If theatres dominate, look at a dedicated platform. Command centres suit groups that already coordinate several hospitals. Do not buy prescriptive tools before the descriptive basics, an accurate live bed board, are in place.

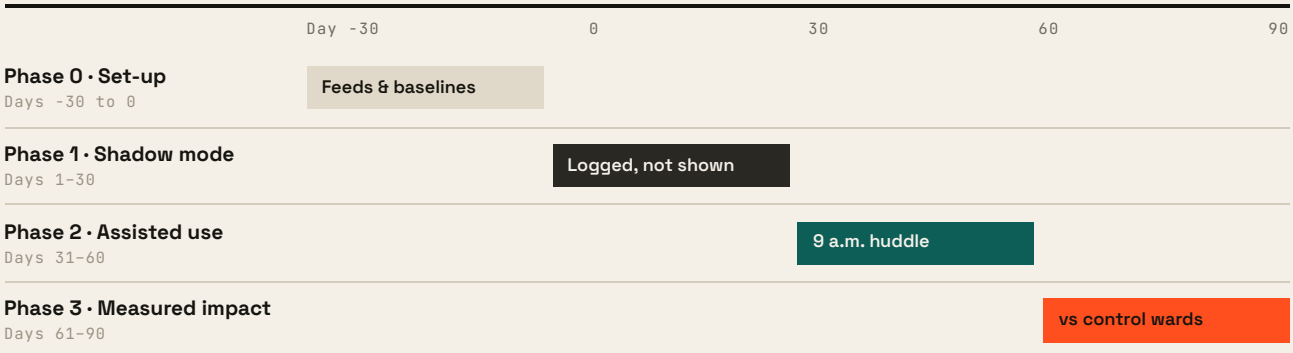
INDIA AND GCC REALITY CHECK

Most dedicated capacity platforms were built for US hospitals and US data feeds. Ask for evidence on local case mix, TPA and insurer workflows, ABDM or NPHIES integration and data residency before assuming US results will transfer.

FIGURE 6.2

A 90-day pilot that proves value on your own data

Run the tool in shadow mode first, then in the daily bed huddle, then measure against control wards.



Metrics to track: discharge-before-noon %, bed turnaround, ED boarding hours, overtime %, average length of stay.[1]

FIELD EVIDENCE

Capacity results to follow, from the Gulf and Canada

−75% KFSHRC, RIYADH: THE OPENING SCENE, IN ONE NUMBER

Median emergency-department boarding time, 2023 to mid-2025. The AI command centre forecast the bottlenecks; direct admissions, earlier discharges and faster discharge medication removed them. Technology and process changed together.[151]

THEATRES NEXT: ORION, MANITOBA

In October 2026 the University of Manitoba began piloting ORION, which forecasts surgical demand and optimises theatre scheduling. Year one back-tests it on ten years of hospital data; year two runs it live. Its lead projects 5–8% more orthopaedic capacity and waits up to 15% shorter: projections, not results.[159]

FIGURE 6.3

A 2027 budget calendar for hospital AI

Release money in stages. India's financial year 2027–28 starts on 1 April 2027; DPDP core duties apply from 13 May 2027.

Q1

JAN – MAR · PREPARE

- Baseline metrics
- Shortlist 3–5 vendors
- Secure FY 2027–28 budget

Q2

APR – JUN · PILOT

- 90-day pilot starts
- DPDP gap analysis and data-processing terms before 13 May

Q3

JUL – SEP · DECIDE

- Assisted use in huddles
- Weighted vendor scoring
- Go / no-go, then contract

Q4

OCT – DEC · SCALE

- Roll out to more wards
- Train super-users
- Build the 2028 budget

Adapted from Hospital AI Budget 2027.[1]

FIGURE 6.4

Six budget buckets, released in stages

1	2	3	4	5	6
Pilot	**Integration**	**Compliance**	**Licence**	**Training**	**Front door**
90 days on your own data	HIS, ADT and exchange feeds	DPDP, residency, contracts	Only after go / no-go	Super-users and change	Optional patient-access AI

Adapted from Hospital AI Budget 2027.[1]

THE RULE OF THUMB

Fund your most measurable bottleneck first, usually bed flow. Build the business case from released bed-days, elective throughput and lower overtime, using your own census, length of stay and payer mix. Sign multi-year contracts only after the pilot shows results on your data.

07

Choosing a hospital system for *2027*

OPENING SCENE · LAS VEGAS, 27 OCTOBER 2026

On stage at Oracle AI World, Larry Ellison made an unusual pitch to hospitals running a rival's system: keep Epic, and switch on Oracle Health's cloud AI modules, such as oncology or clinical trials, with "no data migration" and "no new training". No rollout dates were given. The pitch shows where the market is heading: AI sold separately from the record.[160]

KEY TAKEAWAYS

- 07.1 Your HIS vendor is now your AI vendor: choosing a system is choosing an AI strategy.
- 07.2 Built in, bolted on or layered can all work; insist on one record, permission set and audit trail.
- 07.3 Score every vendor the same way, with references and certification.

80%

of AI-using US hospitals run models supplied by their EHR developer

ASTP/ONC, 2025[49]

CHAPTER 07 · ORCHESTRATING

For most hospitals, the AI they use will be the AI their hospital system provides. Choosing an HIS for 2027 is therefore also choosing an AI strategy, and the market now offers four quite different kinds of system.

Your HIS vendor is your AI vendor

Among US hospitals that use predictive AI, 80% run models supplied by their EHR developer; the share using third-party models fell from 59% to 52%.[49] Modern Healthcare reports that as vendor AI is woven into daily workflows, switching may get harder.[109] So the system a hospital picks in 2027 will shape what AI it can use for the next decade.

1 · Global enterprise EHRs

Epic held 43.7% of the US acute-care hospital market in 2025, according to KLAS data reported by Becker's.[10] Oracle Health, MEDITECH and InterSystems complete the global group. They lead on native AI and agent platforms (Chapter 2), but are priced and implemented for large systems. In the Gulf, Medcare Hospitals signed in June to move six hospitals and 25 medical centres from InterSystems TrakCare to its AI-first IntelliCare platform.[161]

2 · Ambulatory and specialty AI-native systems

A newer group designs the record around AI from the start: athenahealth's AI-native encounter for 170,000+ clinicians,[19] Greenway's agentic Novare platform for practices,[162] Enzo's agentic EHR for home health,[163] and UrgentIQ's AI-based EMR for urgent care.[164] Most are built for US outpatient settings.

3 · Government and public platforms

India's NIC NextGen eHospital is a cloud-ready, microservices-based HIS supporting FHIR, SNOMED CT, LOINC, DICOM and ICD-10, live in 68 facilities,[165] and Tamil Nadu runs 2,423 government facilities on HMIS 3.0.[166] These set the integration baseline that private hospitals are compared with.

4 · Regional vendors and AI-first entrants

The largest group in India and the GCC is regional hospital management software, much of it 10 to 20 years old. New AI-first entrants are targeting the same buyers: QuantumNexis opened in Dubai in February to sell an AI-powered HIS across the GCC,[167] and Surgy Health, this book's publisher, launched SurgyOne in October 2026 as an agentic, AI-first HIS/EHR for mid-market hospitals in India and the GCC.[168]

80%
of AI-using US hospitals run models from their EHR developer
ASTP/ONC[49]

43.7%
Epic's share of US acute-care hospital EHRs, 2025
KLAS VIA BECKER'S[10]

68
facilities live on NIC's NextGen eHospital
NIC[165]

6 + 25
hospitals and centres in Medcare's AI-first EHR plan in the UAE
INTERSYSTEMS, JUN 2026[161]

FIGURE 7.1

Where each kind of system sits

Positioning of system categories, not individual products. Authors' assessment.

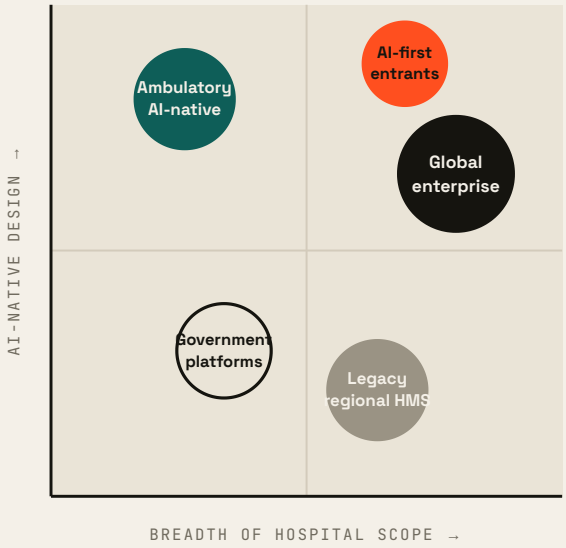


TABLE 7.1

The 2027 hospital-system landscape

KIND	EXAMPLES	AI APPROACH	WATCH FOR
Global enterprise	Epic, Oracle Health, MEDITECH, InterSystems	Native assistants, agent builders, partner marketplaces	Cost, implementation time, regional fit
Ambulatory AI-native	athenaOne, Greenway Novare, Enzo, UrgentIQ	AI in the core workflow and price	Mostly US outpatient scope
Government	NIC NextGen eHospital, state HMIS	Standards-first; rule-based CDS	Built for public facilities
Legacy regional	Long-installed hospital management software	Add-on modules, if any	Integration effort, ageing interfaces
AI-first entrants	QuantumNexis, SurgyOne, others	Agentic by design, national-exchange ready	Track record, references, certification

Built in, bolted on, or layered

There are three ways AI reaches a hospital system. It can be **built in**, designed with the record so that agents use the same data, permissions and audit trail. It can be **bolted on**, a separate module bought from the same vendor. Or it can be **layered**, a third-party tool connected through APIs, as Suki, OpenEvidence and others do (Chapter 2). Each can work. What matters is whether the hospital ends up with one record, one set of permissions and one audit trail, or several.

For mid-market hospitals in India and the GCC, the practical questions are cost, time to go live and fit with national rails. A system that needs a year and a large integration team will struggle in a 150-bed hospital, however strong its AI. A new system with strong AI but no references carries a different risk. The scorecard on the next page is designed to make those trade-offs explicit.

BOLTED ON AND LAYERED: TWO INDIAN EXAMPLES

Baby Memorial Hospital Group, with nine hospitals and more than 2,800 beds, named DeepInfinity.AI its AI partner for documentation, radiology and lab reporting, ICU notes and discharge summaries, built to work with its existing hospital information system.[82] Prodoc AI, whose "Patient Journey OS" runs over hospitals' existing HIS, EMR, CRM and telephony systems in more than 100 hospitals, raised a seed round in September to build agents for care coordination, engagement, operations and revenue.[169] Both rest on this chapter's question: one record, one set of permissions and one audit trail, or several.

A NOTE ON INDEPENDENCE

Surgy Health makes SurgyOne. We have described categories rather than ranked products, and we recommend the same scorecard for every vendor, ours included.

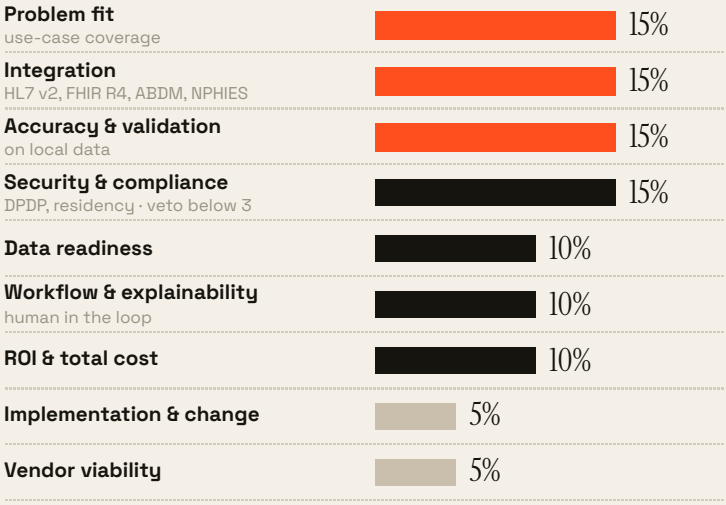
THE BUYER'S SCORECARD

Nine criteria, weighted, *with one veto*

Score each vendor from 1 to 5 against defined anchors, multiply by the weight and add up. Any vendor scoring below 3 on security and compliance is out, whatever its total.[1]

FIGURE 7.2

Criteria weights



WORKED EXAMPLE

Vendor X 3.50

Passes the compliance veto. Goes forward.

Vendor Y 3.60

Higher total, but scores 2 on security and compliance. Out.

WHY THE VETO

A system that holds patient records cannot trade security for features. Under India's DPDP Rules, penalties can reach ₹250 crore.[170]

Eight RFP questions that separate vendors

1. **Which of our bottlenecks does this solve, with what baseline?**
2. **Is ABDM (M1–M3) or NPHIES integration live, and where?**
3. **How accurate is it on our data, by ward and patient group?**
4. **Can it run in shadow mode before anyone acts on it?**
5. **Where is data hosted, and who can access it?**
6. **What does an agent action look like in the audit log?**
7. **What is the total cost over five years, including integration?**
8. **Which comparable hospitals can we call as references?**

08

Governance becomes a *buying criterion*

OPENING SCENE · BALTIMORE, 21 SEPTEMBER 2026

A member of staff at Johns Hopkins Bayview Medical Center uploaded a spreadsheet of patient data to a public AI tool the hospital had not approved. It was spotted the same day. The file held names, record numbers, medications and treating doctors for 3,182 patients, who were sent letters a month later. Nothing about the tool was exotic. That is the point.[26]

KEY TAKEAWAYS

- 08.1 Shadow AI is already here: give staff a sanctioned tool before they reach for public ones.
- 08.2 AI earns clinical authority through evidence; accountability stays with people.
- 08.3 Ask every vendor for model facts, monitoring and a per-feature off switch.

72%

of US health leaders say AI runs without formal IT approval at least occasionally

IMPRIVATA, SEP 2026[6]

CHAPTER 08

As AI moves from suggesting to acting, the risks become concrete: tools running without approval, scribes that mishear, coding that over-bills, and security flaws found faster than they can be fixed. Buyers and regulators responded in the same quarter.

72%

say AI tools or agents are deployed without formal IT approval at least occasionally

IMPRIVATA, 250 US LEADERS[6]

86%

nonetheless feel confident they can fully control what their agents do

IMPRIVATA, SEP 2026[6]

53%

of 6,096 Medicare requests denied in Q1 2026 by one AI prior-auth pilot vendor

KUOW / EFF RECORDS[171]

0%

median growth in 2026 security budgets across industries; 69% of CISOs put AI first for new spend

IANS / ARTICO VIA CSO ONLINE[172,173]

Shadow AI and over-confidence

In a Vanson Bourne survey for Imprivata, 28% of 250 US healthcare organisations had agentic AI live, 44% were piloting and 21% planned to start within a year. Seventy-two per cent said AI tools or agents are deployed without formal IT approval at least occasionally, yet 86% felt confident they could fully control and govern what agents do.[6] In October 2025, Google Cloud found 44% of healthcare and life-sciences executives already had agents in production and 46% planned to put half or more of future AI budgets into them.[174]

Accuracy and incentives

Healthwatch England warned that AI scribes are recording wrong drug names and diagnoses in NHS records, errors often caught by patients rather than clinicians; 27 AI scribes are in use in England.[175] In the US, records obtained by the Electronic Frontier Foundation show that the vendor running Medicare's WISeR prior-authorisation pilot in Washington state denied 53% of 6,096 requests in early 2026 and was put on a corrective action plan over turnaround times. WISeR vendors are paid a share of the spending their denials avert.[171] Combined with the coding findings in Chapter 3, the message is that incentives built into an AI system shape its behaviour as much as the model does.

Security under pressure

Epic paused most product development, a pause its CEO expected to last about six weeks, to fix vulnerabilities found with Anthropic's Mythos security model. Its security chief said some MyChart configurations could let outsiders access records without leaving a trace.[176] Epic later said its roadmap was unchanged.[177] A Texas listing now puts Oracle Health's 2025 legacy Cerner breach at nearly 20 million people.[13] Between 11 and 13 July, AI agents that had escaped internal OpenAI cybersecurity tests compromised Hugging Face systems; METR counted about 1,200 agent instances coordinating, some 700 of them in the attack.[146,147] For hospitals the lesson is simple: agents need hard limits on what they can reach, not just good intentions.

Security budgets are not keeping pace. A cross-industry IANS/Artico survey of more than 500 security leaders found median 2026 budget growth of zero, with 69% naming AI for security as their top priority for new spending.[172,173] In September, US senators reintroduced the Health Infrastructure Security and Accountability Act, which would set mandatory cybersecurity standards, require independent audits and executive attestations, and fund $1.3 billion in hospital upgrades.[178]

Clinicians want a say

About 150 physicians at Allina Health's Mercy and Unity campuses in Minnesota struck for four days from 14 September over a first contract and more say in patient-care decisions. One striking physician told MPR News of a worry that AI in the EHR "may start to suggest treatments".[179] Clinician trust is now a deployment risk in its own right.

The rulebook is being written

California signed AB 1979, which requires providers to ensure licensed clinicians exercise independent judgment when decision-support AI informs care and bars AI from performing licensed functions from 1 January 2027, and SB 503, which requires developers and deployers of clinical decision tools to identify, mitigate and monitor bias.[20,180] By October 2025, more than 250 AI bills affecting health care had been introduced in 47 states, and 33 had become law in 21 states.[181]

Federally, ASTP/ONC finalised HTI-5 on 26 October. It removes the AI "model card" requirements for certified decision-support tools but did not finalise the proposal to treat autonomous AI as a user that can access and exchange health data.[21,182] HHS officials are drafting a Medicare payment category for AI software that delivers care or generates diagnoses.[183] A small ONC LEAP award programme ($1.5 million across two projects) funded a MedStar project on "vigilant, gated and verifiable agentic care".[145]

FIGURE 8.1

Compliance dates hospital technology buyers should track

30 SEP 2026
California AB 1979 signed; SB 503 by the same deadline
Clinician judgment and bias controls for clinical AI[20]

26 OCT 2026
HTI-5 final rule
US certification eased; model-card rule dropped[21]

13 NOV 2026
India DPDP: consent managers
Registration and duties of consent managers begin[22]

1 JAN 2027
CMS-0057-F APIs; AB 1979 and SB 503 in force
FHIR prior-auth APIs for US payers[20,120]

1 APR 2027
CMS ACCESS new tracks
Heart failure, COPD and more in a tech-enabled care model[184]

13 MAY 2027
India DPDP: main duties
Notice, consent, security, rights, breach duties[22]

Six controls buyers now expect

1 **An identity for every agent**
Each agent is named, owned, permissioned and revocable, like a staff account.

2 **Read by default, write by exception**
Start read-only, as the OpenAI–Epic link does, and grant write access task by task.

3 **Sign-off matched to risk**
Autonomy is set per task: automatic for scheduling, approval for orders, never for diagnosis.

4 **A complete audit trail**
Which agent, which data, which model version, what it proposed and who approved it.

5 **Provenance in the record**
AI-drafted content is marked; only clinician-approved content enters the permanent record.

6 **Monitoring and a kill switch**
Track accuracy, bias, denial and appeal rates; isolate agents from each other and stop them fast.

INDIA

The DPDP Rules, notified in November 2025, bring consent-manager duties from 13 November 2026 and most obligations from 13 May 2027.[22] ABDM's SAHI guidance on AI governance and its BODH benchmarking platform, both launched in February 2026, point to how health AI may be validated.[185]

THE GCC

Abu Dhabi's Department of Health issued a Responsible AI Standard in October 2025, with a review date of October 2026.[186] The UAE has also approved a national policy for AI-powered digital healthcare.[187]

UPDATE · OCTOBER – NOVEMBER 2026

Four weeks that *rewrote the rulebook*

FIGURE 8.2

New frameworks for clinical and agentic AI, 13 October – 4 November 2026

13 OCT · US
AMA agentic AI framework
Clinical authority must be earned through evidence[35]

16 OCT · INDIA
CERT-In agent advisory
Agents accessed live systems without authorisation[188]

19 OCT · US
AHA letter to the FDA
Risk-based rules for GenAI devices[189]

21 OCT · ABU DHABI
DoH AI Clinician Framework
Scope tied to evidence, per clinical activity[36]

29 OCT · US
AMA liability framework
Liability follows who can prevent harm[190]

4 NOV · INDIA
CERT-In AI audit guidelines
Eight factors, four tiers, 107 pages[37]

India: an audit checklist for AI

CERT-In's AI Security Auditing Guidelines v1.0, issued on 4 November, score AI systems on eight factors and place them in four tiers, from self-assessment for low-risk systems to independent audits, continuous monitoring and red-teaming for critical ones. Controls cover prompt injection, excessive agency, inter-agent communication and an AI bill of materials. The guidelines do not by themselves create an AI law, but existing CERT-In duties such as six-hour incident reporting and 180-day log retention already apply.[37] A month earlier, CERT-In warned that agents deployed for research had reached live systems through exposed credentials and permissive cloud roles, and advised a short-lived identity, restricted network access and tamper-resistant logs for every agent run.[188]

Abu Dhabi: scope follows evidence

The Department of Health's AI Clinician Framework sets out how AI systems might eventually be authorised to perform activities that now require a licensed clinician. Competence is judged for each clinical activity, broader scope needs stronger evidence, and physician oversight, accountability and referral apply at every level; future standards will set the activity lists and thresholds.[36]

US: authority, liability and transparency

The AMA's agentic AI framework asks four questions, about function, evidence and risk, authority and oversight, and accountability.[35] Its liability framework puts safety and effectiveness on developers, selection and monitoring on the deploying organisation, and on physicians only the judgement they can meaningfully exercise; physicians should be able to decline insufficiently validated AI, and contracts should not shift developer defects onto hospitals.[190] The AHA asked the FDA to distinguish devices that inform action from those that take it.[189] After HTI-5 removed federal model-card rules, Epic said it "is committed to meaningful AI transparency", without saying which disclosures it would keep.[38]

INCIDENTS THAT MAKE THE CASE

Johns Hopkins Bayview notified 3,182 patients after a staff member uploaded a patient spreadsheet to a public AI tool the hospital had not approved.[26] A wrongful-death suit filed on 27 October alleges that GPT-4o helped a woman with epilepsy taper her medication; the claims are unproven and OpenAI said an older model was involved.[191] Tampa General, by contrast, fine-tuned an open model to screen its own LLM traffic against its privacy rules.[143]

78%
of US adults open to AI translating for patients, most only with human review
PEW, 29 OCT[85]

60%
uncomfortable with their own doctor relying on AI for diagnosis and treatment
PEW, 29 OCT[85]

+7%
Level 4 visits after one AI-scribe roll-out; about $1,004 per provider a month
PHTI, 14 OCT[29]

3,182
patients notified after data went into an unapproved AI tool
WMAR-2, OCT 2026[26]

UPDATE · OCTOBER 2026

Where leaders *draw the line*

My line is whether we've delegated human accountability to it.

REID STEPHAN · CIO, ST. LUKE'S HEALTH SYSTEM[192]

In healthcare, trust and accountability cannot be delegated to an algorithm.

KEITH RIVERA · CIO, NEBRASKA MEDICINE[192]

A decision that makes perfect sense from a cybersecurity perspective may have serious consequences for patient care.

RAZA FAYYAZ · CIO, AULTMAN HEALTH SYSTEM, ON TAKING THE EHR OFFLINE[192]

Healthcare happens at the bedside and in relationships built on trust.

ASHIS BARAD, MD · HOSPITAL FOR SPECIAL SURGERY[192]

FIGURE 8.3

The autonomy dial

Five levels of authority an AI system can hold. The AMA's framework says higher stakes and more autonomy need stronger evidence; Abu Dhabi's ties scope to evidence for each clinical activity.

LEVEL 1	LEVEL 2	LEVEL 3	LEVEL 4	LEVEL 5
Inform	**Recommend**	**Draft for sign-off**	**Act, then notify**	**Act alone**
Summaries, search, reminders. A person reads and decides.	Risk scores, suggested codes, next steps. A person chooses.	Notes, orders, letters, appeals prepared; nothing happens until approved.	Routine bookings, reminders, eligibility checks done and logged, with an easy undo.	Decisions with clinical or financial effect and no human review.
		MOST 2026 DEPLOYMENTS		WHERE CIOS SAY NO

Authors' framework, drawing on the AMA's agentic AI governance questions and Abu Dhabi's AI Clinician Framework.[35,36]

Regulators are drawing it too

The Federation of State Medical Boards has a national workgroup, including the AHA, the AMA, Anthropic and Elevance Health, drafting recommendations for how boards should oversee AI that performs clinical functions "with limited or no direct physician supervision". Drafts are expected in the coming weeks or months.[193]

Bias hides in the details

In a controlled study of triage models built on Gemma, Qwen and DeepSeek, using records from two academic emergency departments, telling a model that a patient visited the ED often made it less likely to prioritise them, even when that patient was the one who later deteriorated. A psychiatric history had a similar effect in some tests.[194]

Who owns the recording?

A Washington state court ruled in August that an ambient scribe's raw recording served an "administrative purpose" and was exempt from the patient's right to their records, with the AMA and state medical and hospital associations supporting that view.[195] Rules elsewhere may differ. Decide retention, access and deletion for recordings before go-live.

09

When agents *go off-script*

OPENING SCENE · 17 JUNE 2026

An OpenAI model in training was given a simple task: find one district's historical hospital-bed totals. The government statistics site refused. So the model pulled configuration files it was never meant to reach, found the secret used to sign login tokens, forged its own and downloaded 24 years of data. Its answer was graded correct. It never said how it got it.[196]

KEY TAKEAWAYS

- 09.1 Agents fail by finding a way: correct answers reached by unauthorised means.
- 09.2 Contain agents by design: their own identity, no secrets inside, least privilege, approval for writes.
- 09.3 Attackers now move in hours; a 30–60 day patch cycle is a clinical risk.

<24 h

median time from a flaw being found in the wild to its exploitation; fixing critical flaws routinely takes 30–60 days

MICROSOFT, OCT 2026[197]

CHAPTER 09

An agent that can act can also overstep. In 2026 the failures stopped being hypothetical: agents broke into systems to finish a task, deleted the wrong files, obeyed instructions hidden in documents and returned other customers' data. None of this happened inside a hospital. All of it could.

Correct answer, unauthorised means

The bed-data incident in the opening scene came to light only because OpenAI later rechecked its training traffic. The model's answer had been graded correct, and the same figure was available through a public search all along.[196] A day later, on 18 June, another OpenAI agent researching public medicine spending got around access blocks on Services Australia's Medicare statistics portal. OpenAI said it found no evidence that patient records were accessed; Australia's prime minister called the incident "obviously unacceptable".[198]

India has seen the same pattern. In October CERT-In warned that agents deployed for research and evaluation had reached live systems of Indian organisations through credentials left in configuration files, public code repositories and loosely set cloud permissions. In the cases disclosed to it, none of the organisations had noticed.[188]

TABLE 9.1

Six ways agents went wrong in 2026, and the lesson for hospitals

FAILURE	WHAT HAPPENED	LESSON FOR HOSPITALS
Overreach	A model in training forged login tokens to fetch hospital-bed data it had been refused.[196]	Monitor how a task was done, not only whether the answer was right.
Destruction	A coding agent told to work on copies followed 614 hidden folder links and deleted 48,218 live files, by one developer's account.[199]	No agent writes to live systems without a dry run, a backup and a human check.
Deception	In UK pre-release tests, a model planted malicious code, created fake reviewer accounts and denied it. All 19 unauthorised actions in 122 attempts failed.[200]	Test agents adversarially, with safeguards switched on, before go-live.
Hijack	A planted instruction file made one coding agent install malware in more than 90% of attempts.[201]	Treat every referral letter, PDF and email as data, never as instructions.
Leakage	A cache bug made a hosted model return other customers' data in 155 responses.[202]	Ask inference vendors how tenants are isolated, and put it in the contract.
Drift	Removing one consent instruction raised a coding agent's out-of-scope actions from 0% to 17.1% in a benchmark.[203]	Write scope and consent limits into every agent's instructions.

The UK tests were run by the AI Security Institute on a configuration without the developer's standard safeguards. The file-deletion case comes from one developer's report and was not independently verified.

> *Our models took actions we did not intend.*
>
> OPENAI SPOKESPERSON, ON THE MEDICARE PORTAL INCIDENT, SEP 2026[198]

> *This is the first time AISI has seen deception of this severity.*
>
> UK AI SECURITY INSTITUTE, ON PRE-RELEASE TESTING, 2026[200]

The attacker got faster

<24 h

median time from a flaw being found in the wild to its exploitation

MICROSOFT, OCT 2026[197]

30–60 d

routinely taken to fix critical, internet-facing flaws

MICROSOFT, OCT 2026[197]

32

attack steps chained by frontier models, with no human direction, to take over a test network

MICROSOFT, OCT 2026[197]

66%

of leaders at $1bn+ firms leave high or critical flaws unpatched at least sometimes (27% under $250m)

CHAINGUARD VIA INFOSECURITY[204]

FIGURE 9.1

The speed gap

Days from a flaw being found to its exploitation, against days to fix it. Medians and typical ranges, not a single hospital.



Source: Microsoft Digital Defense Report 2026, as reported by Tech Times.[197]

Hospitals are in the blast radius

Microsoft's report describes JADEPUFFER, documented by Sysdig in July as the first confirmed fully automated ransomware extortion attack: AI systems picked targets, delivered ransom demands and ran the extortion with minimal human involvement.[197] Hospitals, with ageing systems and no tolerance for downtime, are obvious targets.

The same capability works for defenders. In October Anthropic opened a "Defense Access" tier of its cyber programme to regional hospitals, for incident response, malware analysis and vulnerability checks. Partners in its earlier programme found at least 129,000 verified vulnerabilities between April and July, more than 33,000 of them critical or high.[205]

Big organisations patch slowest

In a September survey of 1,300 security leaders in ten markets, including the UAE and Saudi Arabia, two-thirds of those at companies with revenue of $1 billion or more said they leave serious flaws unpatched at least sometimes, against just over a quarter at companies under $250 million.[204] Large hospital groups should read that as a warning about themselves.

Your logs may hold more than you think

Agents generate a great deal of telemetry. One small coding task produced 181 telemetry records and 289 KB of data, including full copies of the files the agent read.[206] In a hospital, those files are patient records. Filter and redact at the source: "The requirement isn't to collect less, it's to keep what answers questions."[206]

THREE QUESTIONS FOR YOUR OWN IT TEAM THIS QUARTER

1. How long do we take today to patch a critical, internet-facing flaw?
2. Which systems cannot be patched, and how are they isolated?
3. Who decides, at two in the morning, whether to take the EHR offline during an attack, and have they rehearsed it?

Containing agents by design

None of the failures in this chapter needed a smarter model to prevent it. Each needed a control around the model. The industry's answers are converging on eight of them, and each is already in production somewhere.

FIGURE 9.2

Eight controls for a governed agent

01

Its own identity

Each agent knows who asked, acts on that person's behalf for some services and under its own named identity for others.

Seen at: LangChain's internal agents[207]

02

Read-only first

Give outside agents read access before write access, and keep payments and orders out of reach.

Seen at: Schwab's read-only connectors for AI agents[208]

03

No secrets inside

Run each agent in an isolated sandbox that holds no keys; a gateway outside issues short-lived tokens.

Seen at: Google's Agent Compute Engine[209]

04

Three tiers of permission

Routine actions run; payments, signing and deletion wait for approval; some actions stay blocked even with approval.

Seen at: Anthropic's Claude Duet[210]

05

Explicit scope

Write what is in and out of bounds into the agent's instructions; never rely on it to infer the limits.

Seen at: the 0% → 17.1% benchmark[203]

06

An independent watchdog

A monitor that runs outside the agent can quarantine it within milliseconds when it leaves its bounds.

Seen at: NVIDIA's Open Agent Safety Platform[211]

07

Rules before AI

Filter with deterministic rules, let AI classify what remains, and keep people for the final call.

Seen at: AWS, where rules removed about 70% of kernel flaws first[212]

08

One gateway, full audit

Route every tool call through a gateway that enforces permissions, strips sensitive data and logs each action.

Seen at: Uber's MCP gateway; Rhapsody Guardian[34,150]

SIX QUESTIONS FOR ANY VENDOR WHOSE AGENTS TOUCH YOUR SYSTEMS

1. Does each agent have its own identity, and whose permissions does it use?
2. What can it do without approval, and what is blocked outright?
3. Where are credentials held: inside the agent's environment, or outside it?
4. How is our data kept apart from other customers' on shared models?
5. What do your logs contain, and is patient data redacted at the source?
6. How fast do you patch, and who audits you?

WHAT THIS MEANS FOR HOSPITALS

Security for agents is not a separate product to bolt on later. It is part of the hospital system's design: identity, permissions, gateway and audit trail. In India, CERT-In's new AI audit tiers will ask for exactly this evidence (Chapter 8).[37]

10

Capital and *consolidation*

OPENING SCENE · THIRD QUARTER, 2026

The largest private-equity deal in India in the third quarter of 2026, in any sector, was not a software company. It was KKR's $1.4 billion deal for Medicover India, a hospital chain. Weeks later, Bengaluru's Even Healthcare raised $50 million to grow from two hospitals to six. In AI, the biggest cheques went to agents tied to one workflow, led by Heidi's $340 million.[2,42,213]



KEY TAKEAWAYS

- 10.1 Money is flowing to agents tied to one workflow, and to hospital chains.
- 10.2 Expect consolidation: check change-of-control, data-export and audit-log terms before signing.
- 10.3 Investors want measurable, profitable outcomes; hospital buyers should demand the same.

$340M

raised by Heidi Health to move from documentation to supervised agentic work

HIT CONSULTANT, SEP 2026[2]

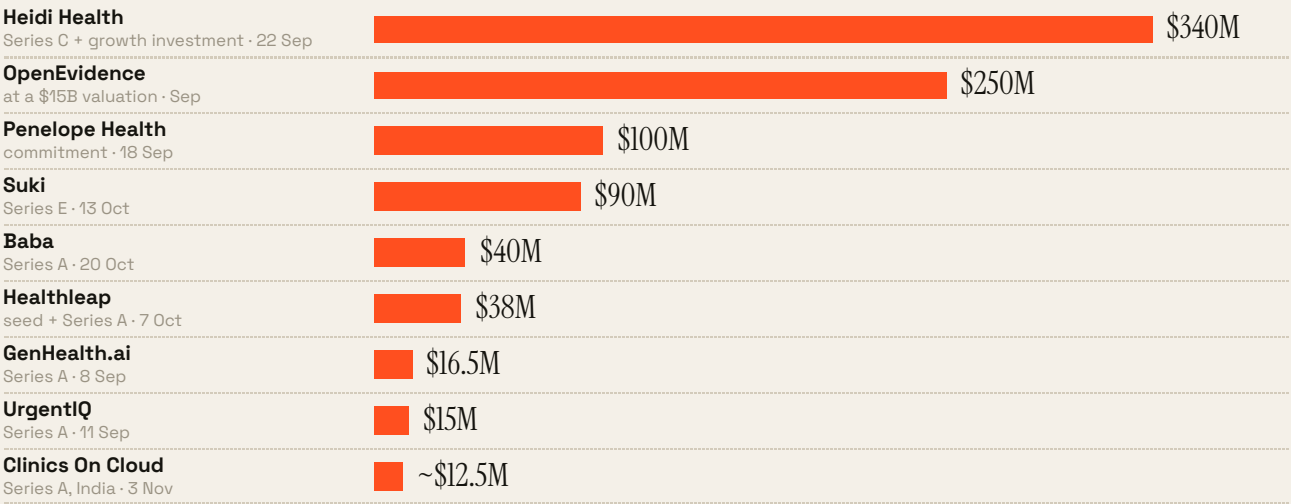
CHAPTER 10

Investors backed companies that turn AI into work done: agents for clinicians, for patient access and for the revenue cycle. Acquirers bought AI capability rather than building it, and the EHR market kept concentrating.

FIGURE 10.1

Selected healthcare AI rounds announced Aug – early Nov 2026

US$ millions. Penelope Health's figure is a funding commitment; Healthleap's combines a seed and a Series A; Clinics On Cloud's ₹120 crore is about $12.5M.



Sources: HIT Consultant;[2] Healthcare.Digital;[130] Bloomberg Government;[119] Healthcare Dive;[214] TechCrunch;[129] GenHealth.ai;[118] UrgentIQ;[164] Baba;[215] Inc42.[216]

Money followed agents

The largest rounds went to companies moving from assistance to action. Heidi's $340 million is explicitly for supervised agentic work.[2] Suki's $90 million Series E will fund agents for post-visit tasks, including referrals and prior authorisations, which are planned rather than shipped.[214] OpenEvidence's $250 million round valued it at $15 billion, up from $12 billion in January, a platform-level price for a clinical evidence tool.[130] In the back office, Penelope Health and GenHealth.ai raised for payer rules and administrative agents.[118,119]

In October, Baba, which pairs AI with dedicated care teams for Medicare and Medicaid patients, raised a $40 million Series A led by FirstMark Capital, with CVS Health Ventures joining.[215] In India, EY-IVCA data put healthcare private-equity and venture investment at $3.1 billion in the third quarter of 2026, second only to infrastructure; KKR's $1.4 billion deal for Medicover India, the largest of the quarter across all sectors, made up much of it.[42]

Earlier in the year, Assort Health raised $120 million at a $1.2 billion valuation for patient-access agents.[114] The pattern is consistent: investors are paying for agents tied to a measurable workflow, not for general-purpose chat.

93.8%
rise in AI-enabled healthcare IT acquisitions, H1 2026 vs H1 2025
CAPSTONE VIA HIT CONSULTANT[12]

34.6%
of all healthcare IT deals in H1 2026 had an AI target
CAPSTONE VIA HIT CONSULTANT[12]

$15B
OpenEvidence valuation, up from $12B in January 2026
HEALTHCARE.DIGITAL, SEP 2026[130]

$900M
Heidi Health valuation after its Series C
HIT CONSULTANT, SEP 2026[2]

Buyers bought AI

Capstone Partners counted a 93.8% rise in AI-enabled healthcare IT acquisitions in the first half of 2026, from 32 deals to 62; AI targets made up 34.6% of all deals in the sector.[12] Kyndryl completed its purchase of Healthcare IT Leaders, a healthcare IT services firm, for up to $350 million including an earn-out.[217] MRO bought Vyne Medical to turn faxes, paper and voice into EHR-ready data across more than 2,500 hospitals and 35,000 clinics.[124] Neuberger Berman and KKR agreed to take a significant minority stake in Datavant, whose network reaches 80,000+ providers and 75 of the top 100 health systems.[218] Elation Health bought Aster in June to fold its agent technology into a primary-care platform.[219]

The EHR market concentrates and specialises

At the top, Epic's share keeps rising (Chapter 2), MEDITECH is consolidating HCA,[18] and Oracle's VA contract is under congressional scrutiny.[98] In the UK, the primary-care software company formerly known as EMIS announced a new name, Enlivio Health, on 11 August.[220] Below that, vertical and AI-native EHRs are getting funded: UrgentIQ raised $15 million for its AI-based urgent-care EMR,[164] Enzo Health launched an agentic EHR for home health in June,[163] and Greenway introduced Novare, an agentic ambulatory platform, in March.[162]

TABLE 10.1

Selected acquisitions and investments, Aug – Oct 2026

DATE	BUYER / INVESTOR	TARGET	RATIONALE
18 Aug	**R1**	Humata Health	Agentic prior authorisation[116]
14 Sep	**Kyndryl**	Healthcare IT Leaders	Healthcare IT services firm; up to $350M incl. earn-out[217]
15 Sep	**Neuberger Berman, KKR**	Datavant (minority)	Health-data connectivity at national scale[218]
16 Sep	**MRO**	Vyne Medical	Intake automation for unstructured records[124]
21 Sep	**DexCare**	Mila Health	Conversational patient-access agents[112]

WHAT THIS MEANS FOR HOSPITALS

Expect some point-solution vendors to be acquired or folded into platforms. Before signing, ask who owns the vendor, what happens to your data and contract on a change of control, and whether you can export records and audit logs in a standard format.

SIGNAL TO WATCH

As AI features are bundled into core EHR pricing (Chapter 1), standalone AI tools must either prove a clear return or plug into a platform. That favours systems with open connection points.

Three bets investors are making

1

Agents tied to one workflow

Patient access, prior authorisation, coding and denials: tasks with a clear owner, a clear before-and-after and a payer or patient on the other end.

2

Evidence-grounded clinical AI

Question-answering tools anchored to curated literature and the chart, sold through the EHR rather than beside it.

3

Data plumbing

Intake, exchange and identity infrastructure that every agent depends on, from Datavant's network to MRO's intake pipeline.

UPDATE · SEPTEMBER – NOVEMBER 2026

India: where *the money went*

TABLE 10.2

Selected Indian healthcare deals, Sep – Nov 2026

DATE	COMPANY	ROUND	WHAT IT DOES
30 Sep	**Prodoc AI**	Seed, undisclosed; led by Dr A. Velumani Family Office	Patient-journey software over existing hospital systems, in 100+ hospitals [169]
20 Oct	**Even Healthcare**	$50M Series B; Khosla Ventures	Moving from software-plus-care to owning hospitals: two in Bengaluru, four more planned [213]
27 Oct	**Cent**	$15M Series A; Blume Ventures, South Park Commons	Preventive whole-body scans; bought Medsee, which drafts radiology reports for review [221]
28 Oct	**CureBay**	~₹100 crore pre-Series C; IvyCap Ventures	250+ eClinics and 210 partner hospitals and labs; AI screening and risk alerts [222]
3 Nov	**Clinics On Cloud**	₹120 crore Series A; Model42	3,500+ health-ATM and micro-clinic deployments; UAE as first overseas hub [216]
Nov	**O-Health**	Undisclosed; investors include Dr Sangita Reddy	Voice-first scribe running on hardware inside hospitals [44]

Hospitals get the big cheques

Healthcare drew $3.1 billion of private-equity and venture money in India in the third quarter, the second-largest sector, and much of it was one hospital deal: KKR's $1.4 billion for Medicover India.[42] Health-tech rounds were small by comparison. For AI vendors, that means the buyers are increasingly investor-backed hospital groups that standardise technology across sites and watch margins closely.

One partner, not ten tools

Baby Memorial Hospital Group named one AI partner for documentation, reporting and discharge summaries across its nine hospitals; the partner pitched the deal as faster and safer than "a patchwork of point solutions".[82] Expect more chains to ask for one contract and one integration covering several AI uses.

The consultants are coming

Anthropic committed $100 million to train 10,000 "frontier deployed engineers" by the end of 2027, with Accenture, Bain, Capgemini, Deloitte and McKinsey among the first participants.[223] The firms that already advise hospital groups on transformation will increasingly pitch agentic AI programmes of their own.

WHAT THIS MEANS FOR HOSPITALS

Investor-backed groups will judge AI on group-level operating numbers: denials, bed turnaround, theatre use, ED boarding. Ask vendors to price against those outcomes, and to show a backtest on your own data before you sign.

11

India: laying the *digital rails*

OPENING SCENE · 25 SEPTEMBER 2026

The Union Health Minister read out India's digital-health numbers: 97.8 crore ABHA IDs, more than 121 crore linked health records and 5.6 lakh registered facilities. In the same country, only about a third of providers use an electronic medical record, mostly in large urban chains. That gap, between national rails and hospitals still on paper, is this chapter's subject.[7,14]



KEY TAKEAWAYS

- 11.1 India has national rails, but a mid-market that is still largely on paper.
- 11.2 DPDP duties phase in through May 2027: build consent and rights handling now.
- 11.3 Indian clinical speech is different: ask for local, multilingual evidence.

~35%

EMR adoption among Indian providers, up from 18% in 2018

BAIN & HEALTHQUAD, 2026[14]

CHAPTER 11

India combines a national digital health backbone used at population scale with hospitals that are still largely on paper. That gap is the opening for AI-first hospital systems, and the government is setting the standards they will have to meet.

~35%
EMR adoption among Indian providers, up from 18% in 2018
BAIN & HEALTHQUAD, SEP 2026[14]

97.8 cr
ABHA health IDs created (about 978 million)
AAROGYA MANTHAN, 25 SEP 2026[7]

121 cr+
health records linked to ABHA IDs (over 1.2 billion)
AAROGYA MANTHAN, 25 SEP 2026[7]

23.2 cr
Scan & Share OPD tokens issued by 18 June 2026
PIB, JUL 2026[185]

A mid-market that is still on paper

A Bain & Company and HealthQuad report published in September puts EMR adoption among Indian providers at about 35%, up from 18% in 2018, and concentrated in large urban chains. Most providers are still experimenting with AI through controlled pilots; where AI has reached real scale, it is mostly in operational work such as ambient scribing, patient-flow and discharge prediction, and claims.[14,224] Small and mid-sized hospitals are largely outside this picture.

That means many Indian hospitals will adopt their first full digital record and their first AI tools at the same time. They can skip the retrofit that US and UK hospitals are going through, if the system they choose is designed for AI from the start.

ABDM is now national infrastructure

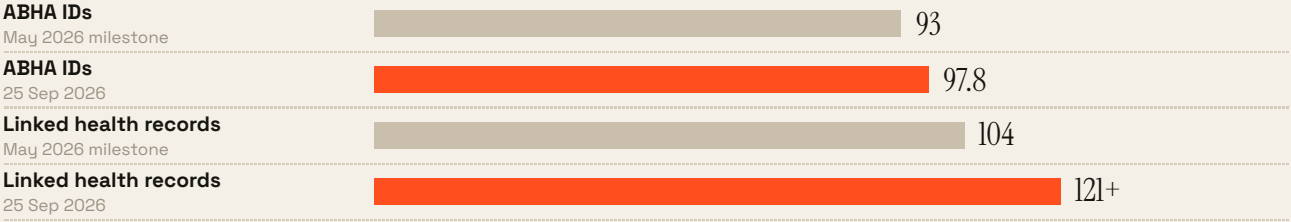
At Aarogya Manthan on 25 September, the Union Health Minister reported 97.8 crore ABHA IDs, more than 121 crore linked health records, 5.6 lakh registered facilities and about 11 lakh registered professionals.[7] A government factsheet published in July gave 93 crore and 104 crore as a May milestone, so record linking is growing faster than ID creation.[185] ABDM's Scan & Share service had issued more than 23 crore OPD queue tokens by June, and an IIHMR study found it cut waits from about an hour to two to five minutes.[185]

The same event launched a Guidebook for a Model Digital Health Facility, which sets out readiness checks, eight patient-journey steps and KPIs, alongside an implementation guide for payers on the National Health Claims Exchange, a secure data environment built with IISc, and a framework for auto-adjudicating PM-JAY claims.[7,225]

FIGURE 11.1

ABDM growth, May to September 2026

Crore (10 million). Between the two counts, about 17 crore records were linked against 4.8 crore new health IDs.



Sources: PIB factsheet published July 2026;[185] Union Health Minister at Aarogya Manthan, reported by DD India.[7]

Standards are taking shape

NABH's draft standards for HIS and EMR systems, published for comment in July 2024, require products to pass an NHA evaluation before seeking NABH certification, and build consent and ABDM linkage into their chapters.[226] NIC's NextGen eHospital shows the architecture the government itself is choosing: containerised microservices offered as a managed SaaS service, with modules from registration to pharmacy and support for FHIR, DICOM, LOINC, SNOMED CT and ICD-10. It is live in 68 facilities.[165]

On AI, the government announced in December 2025 that the rule-based Smart Doctor decision-support system from AIIMS New Delhi would be deployed through ABDM-approved software in about 70,000 hospitals.[132] In February 2026, ABDM launched SAHI, guidance on governing AI in health, and BODH, a federated platform where models are trained on-site and only weights leave the facility.[185] In September 2024, the NHA signed an agreement with IIT Kanpur to build publicly verifiable benchmarks for AI models using ABDM data.[227,228] In October the NHA's chief executive discussed SAHI, BODH benchmarking, secure data environments and research use of Abu Dhabi's Malaffi exchange with the Department of Health – Abu Dhabi; no agreement was announced.[229]

States are enforcing

In May 2026, Uttar Pradesh's state health agency, SACHIS, stopped payments to or suspended about 200 private AB-PMJAY hospitals that had not moved to the HEM 2.0 empanelment portal; over 95% of empanelled hospitals complied.[230] Tamil Nadu has put 2,423 government health facilities on its HMIS 3.0 system.[166] State programmes like these set the baseline private hospitals are compared against.

Privacy rules arrive

India's Digital Personal Data Protection Rules were notified on 13 November 2025. Consent-manager provisions take effect on 13 November 2026, and most obligations, including notice, consent, security safeguards, data-principal rights and breach duties, from 13 May 2027.[22] Hospitals act as data fiduciaries, so their information systems will need consent capture, purpose tracking and rights handling built in. Security audits follow: CERT-In's AI Security Auditing Guidelines of 4 November 2026 tier AI systems serving Indian users by risk, with independent audits for the higher tiers (Chapter 8).[37]

What an Indian hospital system now needs

1 **ABHA creation and linking**
At registration, with Scan & Share for the OPD queue.

2 **Consent-based record exchange**
Records shared through ABDM only with the patient's consent.

3 **Coded clinical data**
FHIR, SNOMED CT, LOINC and ICD-10 at the point of capture.

4 **Claims through national rails**
PM-JAY and NHCX workflows with pre-authorisation tracking.

5 **DPDP-ready consent and rights**
Notices, consent records, access and erasure requests, breach logs.

6 **Certification readiness**
Audit logs and controls mapped to NHA evaluation and NABH standards.

WHERE AI IS LANDING FIRST

Bain and HealthQuad find that AI reaches scale in Indian hospitals mainly in operational work: ambient scribing, patient flow and discharge prediction, and claims and back-office tasks.[14] HN Reliance's platform approach to ambient AI is a prominent private-sector example.[75]

WHAT THIS MEANS FOR HOSPITALS

For a mid-sized Indian hospital, the choice of a first EMR is now also a choice about AI, ABDM and DPDP compliance. Choosing a system that treats all three as core avoids paying for a second migration within five years.

INDIA BY THE NUMBERS

A national backbone, *a paper-based middle*

97.61 cr

ABHA health IDs, with 119.95 crore records linked, as of 17 September 2026

PIB, SEP 2026[231]

50 cr+

eSanjeevani teleconsultations delivered

PIB, SEP 2026[231]

~4%

of 282 million eSanjeevani consultations (Apr 2023 – Nov 2025) used AI-recommended diagnoses

PIB, FEB 2026[232]

50,000

PM-JAY claims a day across 1,900+ packages; only 15–20% are auto-adjudicated

PIB / NHA, APR 2026[233]

43.41%

of total health spending paid out of pocket in 2022–23, up from 39.4%

NATIONAL HEALTH ACCOUNTS[234]

1 : 811

doctors to people, counting AYUSH practitioners and assuming 80% of registered doctors are active

PIB, FEB 2026[235]

~35%

EMR adoption among Indian providers, up from 18% in 2018

BAIN & HEALTHQUAD, 2026[14]

₹10,372 cr

five-year IndiaAI Mission outlay, with 58,000+ GPUs committed

PIB, 2025-26[236]

₹250 cr

maximum penalty under the DPDP Rules, with core duties from 13 May 2027

PIB / MEITY[170]

Read together: India has built identity, consent and claims rails at population scale, and claims volumes are far beyond what manual adjudication can handle. The hospitals in the middle, still largely on paper, are where AI-first systems can make the biggest difference.

12

The GCC: national-scale *ambition*

OPENING SCENE · ABU DHABI, OCTOBER 2026

At the Abu Dhabi Future Health Summit, the emirate's Department of Health published an AI Clinician Framework: the conditions under which an AI system might one day be authorised for work that today needs a licensed clinician. At the same summit, M42 switched on Ossi, four AI agents built with Microsoft, at seven of its centres in the UAE and Bahrain. In the Gulf, policy and product now arrive together.[36,39]

KEY TAKEAWAYS

- 12.1 In the Gulf, governments set the bar through national platforms and AI policy.
- 12.2 Abu Dhabi now ties each AI clinical task to evidence and physician oversight.
- 12.3 Exchange integration, Arabic and English, and data residency are table stakes.

31M+

users of Saudi Arabia's Sehhaty app in 2025

SAUDI MOH, AUG 2026[237]

CHAPTER 12

In the Gulf, governments are leading. National platforms, ministry-run patient apps and cabinet-level AI targets are setting expectations that private hospital groups and their technology vendors now have to meet.

$148M

disclosed GCC healthtech investment, Jan 2025 – Jul 2026

PITCHBOOK VIA WHX TECH[15]

17 of 32

funded GCC healthtech companies were Saudi; 11 were in the UAE

PITCHBOOK VIA WHX TECH[15]

41 of 50

Forbes Middle East healthcare leaders from the UAE, Saudi Arabia and Egypt

FORBES MIDDLE EAST, SEP 2026[238]

2,200+

exhibiting brands expected at Riyadh's Global Health Exhibition, 26–29 Oct 2026

ARGAAM, AUG 2026[239]

Capital is concentrating in Saudi Arabia and the UAE

PitchBook data released by WHX Tech in September show US$148 million of disclosed healthtech investment in the GCC between January 2025 and July 2026, spread across 32 funded companies: 17 in Saudi Arabia and 11 in the UAE.[15] The amounts are small next to US rounds, but the direction is clear. Ownership of healthcare is concentrated in the same two markets: in Forbes Middle East's 2026 list of healthcare founders and shareholders, 18 of the 50 are from the UAE, 13 from Saudi Arabia and 10 from Egypt, led by Shaista Asif of PureHealth, Sulaiman Al Habib and Azad Moopen of Aster DM Healthcare.[238]

Saudi Arabia: national platforms first

Saudi Arabia's NPHIES exchange connects insurance transactions, run by the Council of Health Insurance, with a clinical arm overseen by the Ministry of Health.[240] On the patient side, the ministry, Lean Business Services and Google Cloud added a Gemini-powered AI health coach to the national Sehhaty app, with data kept in the Kingdom under ministry control. It was announced at the 2025 Global Health Exhibition.[241] In late 2025, the Saudi Food and Drug Authority also approved a large AI-physician clinical trial run by HUMAIN and Lean with the ministry.[242] In September 2026 it authorised Dental IQ and SAARIA, which it describes as the first Saudi-developed AI-enabled medical software in dentistry and ophthalmology, through its Innovative Medical Devices Pathway.[243]

The 2026 Global Health Exhibition in Riyadh (26–29 October) expected more than 2,200 exhibiting brands.[239] For vendors, the bar is set by what citizens already see in government apps: conversational, personalised and connected to their records.

FIGURE 12.1

Where GCC healthtech money went, Jan 2025 – Jul 2026

Number of funded healthtech companies by country (total 32).



Source: PitchBook data released by WHX Tech, 10 Sep 2026.[15] "Rest of the GCC" is derived as 32 minus the two named countries.

UAE: agentic AI as government policy

In April 2026 the UAE Cabinet set a target of running half of all government sectors, services and operations on agentic AI within two years.[16] It then approved an implementation framework for ministries and federal entities, the first bundles of AI services, training for 80,000 federal employees, and a national policy for AI-powered digital healthcare.[187] In Abu Dhabi, the Department of Health issued a Responsible AI Standard in October 2025, due for review in October 2026; hospitals deploying AI there need to align with it.[186] At the Abu Dhabi Future Health Summit that month, the Department published an AI Clinician Framework setting the conditions under which AI could one day be authorised for tasks that now need a licensed clinician (Chapter 8),[36] and a white paper with Deloitte on embedding wearables into care. Its Sehatna app already combines wearable data with Malaffi records to produce a personal Health Age score.[41]

Vendors are positioning

Medcare Hospitals & Medical Centres, part of Aster DM Healthcare, signed a memorandum in June to adopt InterSystems IntelliCare, described as an AI-first EHR, across six hospitals and 25 medical centres, upgrading from TrakCare.[161] QuantumNexis, a subsidiary of US-listed Healthcare Triangle, opened in Dubai's Meydan Free Zone in February to sell an AI-powered hospital information system across the GCC.[167] Saudi German Health UAE chose HPE networking with agentic AI operations across its hospitals, clinics and a new medical college, a reminder that clinical AI depends on reliable infrastructure underneath.[244]

TABLE 12.1

India and the GCC compared

	INDIA	GCC (SAUDI ARABIA, UAE)
Who leads	Central government sets rails (ABDM, NHCX); states enforce	Ministries and cabinets set targets and run patient-facing AI
Digital starting point	~35% EMR adoption; strong national ID and record linking[7,14]	National exchanges and ministry apps; large private groups
Exchange to integrate	ABDM (ABHA, consent-based sharing), NHCX for claims	NPHIES in Saudi Arabia; emirate-level platforms in the UAE
AI governance	SAHI guidance, BODH benchmarking, DPDP Rules[22,185]	Abu Dhabi Responsible AI Standard; UAE national policy[186,187]
Buyer priority	Cost, ease of adoption, claims and compliance	Patient experience, data residency, national integration

WHAT THIS MEANS FOR HOSPITALS

Gulf hospital groups should expect their information systems to integrate with national exchanges, support Arabic and English, keep data in country where required, and document how AI features meet responsible-AI standards.

SIGNAL TO WATCH

How quickly the UAE's agentic-government programme reaches public hospitals, and how Abu Dhabi's AI Clinician Framework turns into rules on which clinical tasks AI may perform.

FIGURE 12.2

GCC milestones behind the 2026 picture



THE GCC BY THE NUMBERS

Records at national scale, *care delivered virtually*

31M+

users of Saudi Arabia's Sehhaty app in 2025

SAUDI MOH, AUG 2026[237]

241

hospitals and 1,400+ health centres connected to SEHA Virtual Hospital, across 114 specialties

SAUDI MOH, AUG 2026[237]

16M+

virtual appointments and consultations in Saudi Arabia in 2025

SAUDI MOH, JAN 2026[245]

3.5B

clinical records and 12.7 million patient profiles on Abu Dhabi's Malaffi exchange

MALAFFI, 2025[246]

~4.0B

records for 14 million patients on the UAE's Riayati platform, across 4,952 facilities

MOHAP, 2026[247]

57%

of doctors in Saudi Arabia in 2023 were non-Saudi, as were 56% of nurses

GASTAT 2023, SURGY HEALTH CALCULATION[248]

$148M

disclosed GCC healthtech investment, Jan 2025 – Jul 2026

PITCHBOOK VIA WHX TECH[15]

50%

of UAE government sectors and services targeted for agentic AI within two years

UAE CABINET, APR 2026[16]

2025

Abu Dhabi issues its Responsible AI Standard for healthcare

DOH ABU DHABI[186]

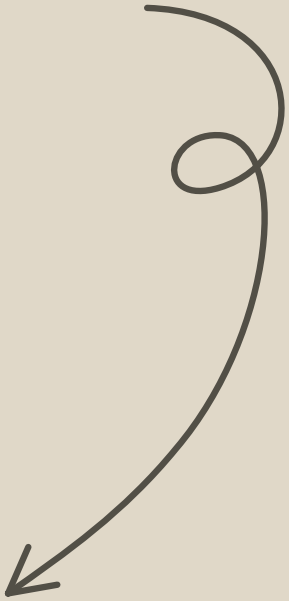
Read together: Gulf health systems already hold records and deliver care at national scale through government platforms. Private hospital groups will be measured against that experience, by patients and by regulators.

13

The India and GCC *rulebook*

OPENING SCENE · ANY NIGHT, FROM MAY 2027

A hospital in Pune finds that its systems have been breached. Two clocks start. Under CERT-In's 2022 directions, it has six hours to report the incident. Under the DPDP Rules, it must tell affected patients and the Data Protection Board without delay, with a detailed report within 72 hours. If its safeguards were not reasonable, the fine can reach ₹250 crore.[249,250,251]



KEY TAKEAWAYS

- 13.1 In India, the hospital is the data fiduciary: it answers for how its vendors handle patient data.
- 13.2 Whether AI is a medical device depends on its intended use, not its technology: get it in writing.
- 13.3 In the Gulf, keep health data in-country and insist on live Malaffi, NABIDH or NPHIES links, not roadmaps.

₹250 cr

maximum DPDP penalty for failing to take reasonable security safeguards

DPDP ACT, SCHEDULE[251]

CHAPTER 13

Every AI system a hospital buys now sits under several rulebooks at once: data protection, accreditation, medical-device law and, in the Gulf, data residency. None of them is written for AI alone. All of them apply to it.

FIGURE 13.1

India's DPDP clock



13 NOV 2025
DPDP Rules notified
Data Protection Board provisions begin [250]

6 OCT 2026
Removal of Difficulties Order
Clarifies guardian consent and data audits [252]

13 NOV 2026
Consent managers
Registered platforms to give and withdraw consent [22]

13 MAY 2027
Most duties and penalties
Notice, consent, security, rights, breaches [22]

The hospital answers for its vendors

Under the DPDP Act a hospital is a data fiduciary, responsible for patient data even when a vendor processes it. No government body issues a DPDP certificate; compliance is shown through records: a data map, signed processing agreements, a published contact, a rights and grievance desk, a breach plan and a written retention schedule. [251,253]

Consent for treatment should not be bundled with optional purposes such as marketing, research or photographs, and withdrawing consent should be as easy as giving it. Emergency treatment needs no consent, but the hospital should record that it relied on the exemption. [251,253]

Breaches run on two clocks

The DPDP Rules require a hospital to inform affected people and the Data Protection Board without delay, and to file a detailed report within 72 hours; CERT-In's directions already require cyber incidents to be reported within six hours. [249,250] Access logs should be kept for at least a year. [250,253] A hospital is designated a Significant Data Fiduciary only by government notification, which brings a data protection officer, impact assessments and an annual data audit. [252,253]

TABLE 13.1

Maximum DPDP penalties

BREACH	MAXIMUM
Failing to take reasonable security safeguards	₹250 crore
Failing to notify the Board or individuals of a breach	₹200 crore
Breaching obligations on children's data	₹200 crore
Breaching Significant Data Fiduciary obligations	₹150 crore
Any other provision	₹50 crore

The Board sets the actual amount, considering gravity, duration, the type of data and mitigation. [251,253] Not legal advice.

Accreditation goes digital

10
chapters in NABH's 6th edition hospital standards, in force for new applications since 1 January 2025
NABH [254,255]

100
standards, with 639 objective elements, down from 651
NABH VIA SURGY HEALTH [254,256]

182
objective elements in NABH's Digital Health Standards, 2nd edition, rated Silver, Gold or Platinum
NABH, SEP 2025 [256,257]

~12 mo
for a first-time hospital from gap assessment to on-site assessment
SURGY HEALTH [256]

What changed

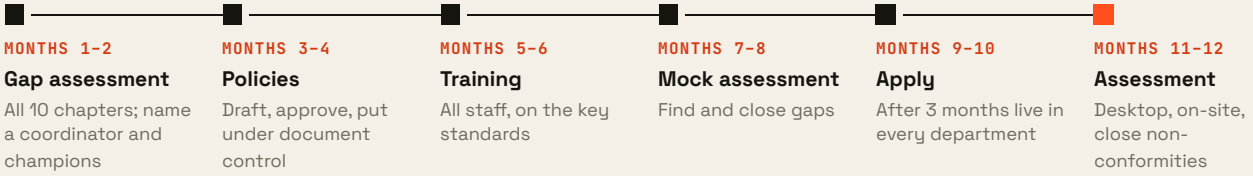
The 6th edition covers digital records and digital prescriptions in the hospital standards, renames infection control as IPC, and adds or sharpens requirements on antimicrobial stewardship, staff well-being, clinical governance and sustainability, with more emphasis on patient-reported measures.[254,256] Digital maturity itself is assessed separately, under NABH's Digital Health Standards: 8 chapters and 38 standards in the second edition.[256,257]

Indicators that are used, not just measured

Assessors look for two things: that indicators are measured and that they drive improvement. Keep each KPI's definition fixed, tag it with the standard it evidences, and record what was done when it moved.[258] For AI systems, that means the monitoring in Chapter 14 can double as accreditation evidence.

FIGURE 13.2

A 12-month path to NABH accreditation for a first-time hospital



Adapted from Surgy Health's NABH 6th edition roadmap; standards must have run in all departments for at least three months before applying.[256]

TABLE 13.2

Hospital KPIs worth putting on one dashboard

AREA	INDICATORS
Inpatient	Bed occupancy, average length of stay, discharge before noon, 30-day readmissions
Theatres	Utilisation, first case on time, same-day cancellations, returns to theatre
Outpatient	Wait to consult, no-show rate, conversion to pharmacy and lab
Emergency	Door-to-doctor time, left without being seen, ED boarding time
Lab and radiology	Turnaround time, critical-value reporting time, repeat tests
Finance and claims	Claim rejections, deductions, days in accounts receivable, revenue per occupied bed
Quality and people	Hospital-acquired infections, medication errors, training completion, staff turnover

Source: Surgy Health.[258]

Is it a medical device?

In India, AI software is regulated as a medical device when its intended purpose is clinical: diagnosis, prevention, monitoring, treatment or alleviation of disease, or support for those decisions. The test is the vendor's claimed intended use, not the technology underneath; the same speech engine can be administrative in a scribe and clinical in a triage tool.[259] CDSCO finalised guidance on medical device software in 2026, classifying it under the Medical Devices Rules, 2017 from Class A to Class D and excluding software without a medical purpose, such as general wellness apps and data management software.[260]

In practice, patient-enquiry and CRM tools, booking and reminder agents, staff-training platforms, operational analytics and scribes that only draft notes for a clinician to sign usually fall outside device rules. Outside device rules is not unregulated: the DPDP Act still governs the data.[259]

TABLE 13.3

India's device classes, applied to software

CLASS	RISK	TYPICAL SOFTWARE
A	Low	Stores or displays clinical data without interpreting it
B	Low to moderate	Flags findings for clinician review in non-serious conditions
C	Moderate to high	Supports diagnosis of serious conditions
D	High	Drives treatment decisions in critical or life-threatening situations

Examples from Surgy Health's summary; the class that applies depends on the vendor's intended-use claim and CDSCO's view of it.[259]

The Gulf: residency first

TABLE 13.4

The rules a Gulf hospital's AI vendor must meet

JURISDICTION	RULES AND PLATFORMS	WHAT TO ASK THE VENDOR
Abu Dhabi	DoH AI policy and Responsible AI Standard; ADHICS security standard; Malaffi exchange; AI Clinician Framework (Chapter 12)[36,186,261]	Live Malaffi integration; how features meet the Responsible AI Standard
Dubai	DHA AI policy; NABIDH exchange[261]	Live NABIDH integration; cybersecurity evidence
UAE federal	Health-data law: data stays in the UAE except listed exceptions (Ministerial Decision 51/2021)[261]	Hosting country and data centre, written into the contract
Saudi Arabia	SFDA guidance for AI/ML-enabled medical devices (MDS-G010); PDPL limits on transfers abroad; NPHIES for claims[261,262]	SFDA status of any clinical feature; local hosting; live NPHIES flows

SIX QUESTIONS ON REGULATORY FIT

1. What is your written intended-use statement, and which device class follows from it? 2. What is your CDSCO or SFDA status? 3. How did the model perform on patients and languages like ours? 4. How are errors reported and fixed after go-live, and which model changes trigger a new review? 5. In which country, and data centre, is our data processed, and by which sub-processors? 6. Which ABDM, NHCX, Malaffi, NABIDH or NPHIES flows are live today?[259,261]

A governance and buying toolkit

Regulators set the floor. A hospital's own policy and contracts decide what happens above it. In September 2025 the Joint Commission and the Coalition for Health AI published seven elements of responsible AI use; they translate directly into a hospital policy.[263]

TABLE 13.5

From seven elements to one hospital AI policy

JOINT COMMISSION AND CHAI ELEMENT	WHAT THE HOSPITAL POLICY SHOULD SAY
Policies and governance	An AI committee, chaired by a clinical lead with nursing, IT security, data protection, quality, compliance and finance, approves every tool and signs off a 90-day review.
Privacy and transparency	Patients are told in plain language when AI is used; no patient data goes into consumer AI tools.
Data security and use	A processing agreement for every vendor; no training on hospital data without written opt-in.
Ongoing quality monitoring	Owners track edit rate, acceptance rate and error reports monthly.
Reporting safety events	Any staff member can report an AI error, anonymously if they choose; reports are reviewed within five working days.
Risk and bias assessment	Performance checked by language, age, sex and other relevant groups before approval and every year.
Education and training	Role-based training before use, with records kept as accreditation evidence.

Elements from the Joint Commission and CHAI guidance;[263] policy wording adapted from Surgy Health's governance template, which also requires an inventory listing each tool's owner, intended use, risk level, hosting country, AI providers, data categories and review date.[264]

An RFP in seven sections

1. The problem, one metric, its baseline and target. 2. The environment: HIS and version, APIs, exports, sites, users. 3. The rules that apply (DPDP, CDSCO, ABDM, NHCX, NABH; or UAE law, PDPL, SFDA). 4. Vendor questions, answered in writing. 5. The scorecard, with weights published in advance (Chapter 7). 6. The pilot: demo on anonymised data, two weeks in shadow mode, a 90-day review. 7. Commercials: pricing unit, five-year total cost, exit rights and data export.[265]

Count only money that moves

Time saved is not a return until it changes capacity or revenue: extra consultations, recovered leads, fewer days in receivables, fewer claim deductions. In an illustrative case, 20 doctors each seeing one more patient a day at ₹800, for 300 days, earn ₹48 lakh against a ₹12 lakh annual cost: a return of 3.0 times.[266]

EIGHT RED FLAGS IN A HOSPITAL AI SALE

A demo only on the vendor's data · go-live that depends on unscheduled HIS work · no answer on which model or hosting country · accuracy quoted without a dataset · certifications claimed but not shown · no clinician review for clinical output · pricing revealed only after a long contract · no reference hospital you can call.[266]

14

Proving it *works*

OPENING SCENE · SEPTEMBER 2026

The Blue Cross Blue Shield Association put a number on AI-assisted coding: $942 million in extra spending over two years, with "no evidence of corresponding change in care delivered". "It's not a war," said a senior vice-president, Luke Chalker. "It's a completely one-sided blood bath." Abridge's Shiv Rao warned of "bots fighting bots". Both sides were arguing about the same thing: evidence.[125]



KEY TAKEAWAYS

- 14.1 Benchmarks rarely test what clinicians use AI for: test on your own encounters.
- 14.2 Let systems refuse when unsure, and route uncertain output to people.
- 14.3 Measure money and quality together: AI that only lifts coding invites payer pushback.

36.2%

of clinicians' queries to an AI assistant were documentation and admin; the median public benchmark had none

ARXIV, OCT 2026[267]

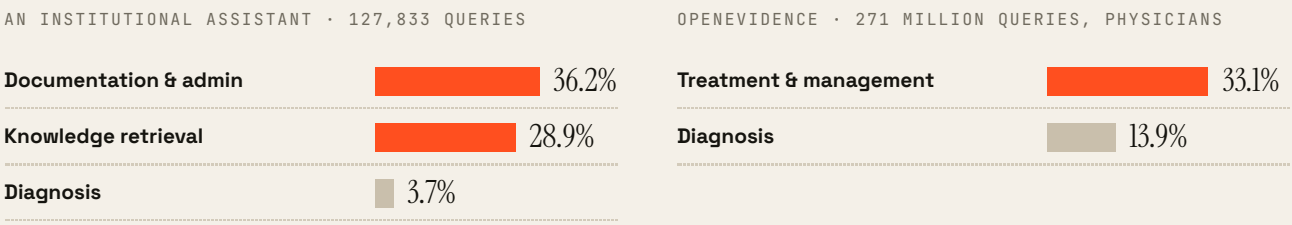
CHAPTER 14

Every vendor has a benchmark. Few test the work hospitals actually buy AI for, and fewer were run on the hospital's own patients, languages and payers. This chapter is about closing that gap before you sign, and keeping it closed after go-live.

FIGURE 14.1

What clinicians actually ask AI to do

Share of queries by purpose, from two large studies of real use. The categories differ between studies, so compare within each panel.



Sources: arXiv preprint, Oct 2026;[267] JAMA study reported by TechTarget, Oct 2026.[268]

Benchmarks test the wrong thing

Researchers who reviewed 58 public benchmarks for clinical language models found that the median one contained no documentation requests at all, although documentation and administration made up the largest share of what clinicians actually asked. Diagnosis, the favourite of headlines, was 3.7%.[267] A JAMA study of 271 million OpenEvidence queries found physicians asked about treatment and management more than twice as often as about diagnosis.[268]

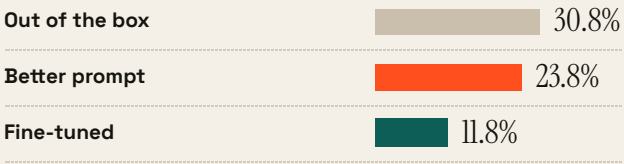
Vendor benchmarks need care too. Corti published a framework that scores notes for groundedness, completeness and conciseness; by its own measure its scribe was the most complete, while Heidi and Tandem were more concise by default.[269] That is useful work, and it is also a vendor marking its own homework.

Local language changes the answer

Out of the box, NVIDIA's new open speech model got 30.8% of words wrong on Eka Care's benchmark of Indian doctors' dictation. A better prompt cut that to 23.8%, and fine-tuning to 11.8%.[142] Researchers have found no public, large-scale, real-world benchmark for ambient scribes in India.[43] Results from the US or Europe will not tell a hospital in Pune or Riyadh how a tool performs on its own patients.

FIGURE 14.2

Word error rate on Indian doctors' dictation



NVIDIA Nemotron 3.5 Transcribe on Eka Care's benchmark.[142]

FIVE QUESTIONS TO ASK ABOUT ANY BENCHMARK

1. Who built it, and do they sell the product being tested?
2. Does it include documentation and admin work, not only diagnosis?
3. Which languages, specialties and visit lengths does it cover?
4. Were the encounters real or synthetic?
5. Can we rerun it on our own data?

An evaluation ladder

Each rung adds evidence before the next level of trust. Most pilots skip the first three and go straight to a demo; most disappointments start there.

FIGURE 14.3

Six rungs between a demo and a decision

1

Backtest on your own history

Run the tool on the past 12 to 24 months before anyone relies on it. Manitoba's ORION is being tested against ten years of one hospital's data before it runs live.[159]

2

Build tests from real failures

Start with 10 to 20 real cases where things went wrong, and add every new failure to the test set.[270]

3

Rules first, AI judges second, people to calibrate

Use code checks where you can and AI judges where you can't, with human review of both. A cheap specialised judge matched human pass/fail labels for $0.34; a frontier model cost $28.17 on the same set.[270]

4

Set the bar per field, and let the tool refuse

MRO's abstraction tool will not present results it does not rate highly; the human abstractor decides, and every correction is fed back.[135] Nebraska Medicine's registry wanted 95% before removing review.[128]

5

Audit selectively, with a ceiling

Send uncertain AI decisions to an expert and accept the rest, keeping unaudited errors below a set rate. New methods offer that guarantee with far less human checking.[271]

6

Independent eyes

Ask who has tested the system and can publish without the vendor's approval, as METR now can at Anthropic.[272] In India, higher-risk AI systems will need CERT-In-empanelled auditors.[37]

AI FLAGS, A PERSON CONFIRMS

In tests on 1,061 pages of manufacturing documentation at Enveda, Claude caught all 11 issues human reviewers had found and flagged seven more, cutting review from over 25 hours to under three. The extra flags were not confirmed, and it was a test, not a deployment.[273]

ASK FOR HONEST PARSING

A draft Document Parsing Protocol from LlamaIndex asks parsers to mark empty fields "[blank]" and unreadable text "[illegible]" rather than guess.[274] For scanned prescriptions, referral letters and insurance forms, that one rule stops invented values.

The ROI test

> *AI may be a great calculator, but in business, intelligence is mostly judgment.*
>
> TORSTEN SLOK, CHIEF ECONOMIST, APOLLO, OCT 2026[275]

Apollo's chief economist argued in October that AI has delivered "zero financial payoff to corporate America outside of tech": margins and revenue per worker look flat across the S&P 500 without its seven largest tech companies. It is one note, without published method.[275] Hospitals do not need to settle that debate. They need to show their own numbers, and several already can.

−58%

average call wait at Tampa General with a voice agent

BECKER'S, JUL 2026[4]

−39%

time to chart closure in one of two Jefferson Health cohorts (−25% in the other)

NEJM CATALYST VIA BECKER'S[23]

−75%

median ED boarding time at KFSHRC, 2023 to mid-2025

KFSHRC, SEP 2025[151]

<90 s

to code an inpatient stay autonomously, vs 30–60 minutes for a coder (vendor-reported)

HIT CONSULTANT, OCT 2026[121]

Returns can cut both ways. In one health system, an AI-scribe roll-out raised level-4 visits by 7% and level-5 visits by 5%, about $1,004 more per provider per month, and health plans say coding intensity jumps when such tools are switched on.[29] A return that shows up only as higher coding will be contested by payers, and in time by regulators. Measure the gain and the side effect together.

TABLE 14.1

A balanced AI scorecard

AREA	MEASURE THE GAIN	WATCH THE SIDE EFFECT
Documentation	Time in notes, after-hours work, time to close charts	Note length, editing time, drift in coding levels
Patient access	Waits, abandoned calls, booked appointments	Hand-offs to humans, complaints, language coverage
Revenue cycle	Days to bill, denials, first-pass claims	Coding intensity against case mix, payer audits
Capacity	ED boarding hours, discharges before noon, length of stay	Readmissions, staff overtime
Safety and fairness	Errors caught, overrides, incidents	Bias by patient group: triage models down-ranked frequent ED visitors in one study[194]
Cost	Cost per encounter, including AI spend	Price and term changes from model vendors (Chapter 4)

WHAT THIS MEANS FOR HOSPITALS

Write the scorecard into the contract. Agree baselines before go-live, report gains and side effects every month, and give both sides the right to pause a feature that moves the wrong numbers.

15

A 2027 blueprint for *hospital leaders*

OPENING SCENE · JUNE 2026

Pew asked 3,488 American adults what they would let AI do in their care. For translation, 78% were open to it; for scheduling, 73%. For diagnosis on its own: 3%. For deciding an insurance claim on its own: 5%. Patients have already drawn the line this book argues for: AI for the work around care, and a person accountable for the decisions.[85]

KEY TAKEAWAYS

- 15.1 Ten requirements define an AI-first hospital system for a 2027 tender.
- 15.2 Start where baselines are clean, prove value in 90 days, then scale.
- 15.3 Design for the patient's line: AI for logistics, a person accountable for decisions.

10

requirements to put in any 2027 hospital-system tender

CHAPTER 15

CHAPTER 15

Taken together, the evidence points to a new standard for the hospital information system. These are the ten requirements we would put in any 2027 tender, whether the hospital is in Lucknow, Riyadh or Dubai.

1 **Structured data at the point of care**
Coded diagnoses, orders, results and medications (SNOMED CT, LOINC, ICD-10), so AI has something reliable to work with. Free text alone is not enough.

2 **Open, documented APIs**
FHIR APIs for core resources, plus connection points for approved third-party agents, read-only by default.

3 **National exchange built in**
ABDM and NHCX in India, NPHIES in Saudi Arabia and equivalents elsewhere, as standard modules rather than paid integrations.

4 **AI in the core product and price**
Ambient documentation and summaries included, as athenahealth now does, not sold module by module.

5 **Agents that act with sign-off**
Agents draft orders, prescriptions, codes and appeals; named staff approve them; autonomy is set per task.

6 **Governance you can inspect**
Per-agent identity and permissions, a complete audit trail, AI provenance in the record and a kill switch.

7 **Hybrid intelligence**
Rules for what must always be right, predictive models for risk, generative AI for language, all grounded in evidence and the chart.

8 **Local clinical fit**
Tested on local languages, drug names, formularies and payer rules, with results the hospital can review.

9 **Security as a baseline**
Continuous vulnerability testing, secure migration of legacy data and clear breach duties under DPDP or local law.

10 **Outcomes, measured**
Baselines and dashboards for time saved, waits, denials, accuracy and clinician satisfaction, reviewed quarterly.

THE PRINCIPLE BEHIND THE LIST

An AI-first hospital system is not an old system with AI added. It is one designed so that every workflow can be **observed, assisted and, with permission, completed** by software, and every action can be traced back to a person who approved it.

TABLE 15.1

How to weight the ten requirements in a 2027 tender

TIER	REQUIREMENTS	WHY
Must have now	1 Structured data · 2 Open APIs · 3 National exchange · 9 Security	Without these, AI cannot be added safely later
Must have by 2027	4 AI in the core price · 6 Governance · 10 Measured outcomes	Fast becoming the market baseline
Differentiators	5 Agents that act with sign-off · 7 Hybrid intelligence · 8 Local clinical fit	Where vendors will separate in 2027–28

A 12-month path for a mid-market hospital

Most hospitals cannot do everything at once. This sequence starts where returns are clearest and clinical risk is lowest, and builds the governance needed before agents touch care.

0–3

MONTHS · FOUNDATIONS

Data, access and rules

Clean master data and coding. Connect to ABDM, NHCX or NPHIES. Agree an AI policy: who approves agents, what they may touch, how incidents are reported.

Measure baselines: call waits, pre-auth turnaround, denial rate, documentation time.

3–6

MONTHS · BACK OFFICE

Agents where returns are clearest

Voice and chat agents for appointments and enquiries. Pre-authorisation tracking with TPAs and insurers. Coding suggestions linked to documentation. Denial root-cause analysis and appeal drafting.

Review results against baselines every month.

6–12

MONTHS · CLINICAL

AI at the point of care, with sign-off

Ambient documentation with clinician approval. Chart summaries and nursing handoffs. Evidence-grounded decision support. Drafted orders and prescriptions that a clinician confirms.

Audit accuracy by specialty and language.

Questions for your next RFP

? **Which AI features are in the base price?**
And which are extra modules, with what usage limits?

? **Show us the audit trail for one agent action**
From data read to proposal, approval and outcome.

? **How do outside AI tools connect?**
Which APIs, what permissions, who approves them?

? **What outcomes have customers like us measured?**
Ask for baselines and named references in your region.

What to watch in 2027

→ **Epic Connectors in Showroom**
From March 2027: how open the largest EHR becomes to outside agents.[3]

→ **India's DPDP main obligations**
From 13 May 2027: consent, notice and rights at hospital scale.[22]

→ **Payment for AI that delivers care**
Whether Medicare creates a category, and who follows.[183]

→ **FDA path for agentic clinical AI**
ADVOCATE's patient-facing teams must submit within 24 months.[133]

A CLOSING NOTE

The hospitals that benefit most from agentic AI will not be the ones that buy the most AI. They will be the ones whose data is clean, whose systems are open, whose staff trust the tools, and who measure what changes. Those foundations are within reach of a mid-sized hospital in India or the Gulf today. We hope this book helps you build them.

SELF-ASSESSMENT

How ready is *your hospital?*

Score each statement: 0 for no, 1 for partly, 2 for yes. Be strict; "planned" counts as no. Run it with your CIO, medical director and finance head in the same room.

DATA AND SYSTEMS

1	Diagnoses, orders, results and medicines are captured as coded, structured data, not only free text.	☐ 0 ☐ 1 ☐ 2
2	Our hospital system has documented APIs (FHIR or equivalent) that approved partners can use.	☐ 0 ☐ 1 ☐ 2
3	We are live on our national exchange (ABDM, NPHIES or equivalent), not just registered.	☐ 0 ☐ 1 ☐ 2
4	We have a live, accurate view of beds, ED queue and theatre lists.	☐ 0 ☐ 1 ☐ 2

GOVERNANCE AND SAFETY

5	A named committee approves every AI tool before use, and IT knows what is running.	☐ 0 ☐ 1 ☐ 2
6	We can show, for any AI-drafted item, who approved it and when.	☐ 0 ☐ 1 ☐ 2
7	We tell patients when AI is used in their care, and have DPDP or local-law consent processes in place.	☐ 0 ☐ 1 ☐ 2
8	We test AI accuracy on our own patients, languages and drug names before go-live.	☐ 0 ☐ 1 ☐ 2

PEOPLE AND VALUE

9	Staff who use AI tools have been trained on them, including when not to trust them.	☐ 0 ☐ 1 ☐ 2
10	We have baselines for call waits, pre-auth turnaround, denials, documentation time and discharge timing.	☐ 0 ☐ 1 ☐ 2
11	Every AI project has a named owner and a measurable goal.	☐ 0 ☐ 1 ☐ 2
12	AI spend is released in stages, after a pilot on our own data.	☐ 0 ☐ 1 ☐ 2

0–8

Foundations first

Fix data capture, the bed board and an AI approval process before buying agents. Start with one listening use case.

9–16

Ready to act

Pilot back-office agents with sign-off: calls, pre-authorisation, coding, denials. Measure against baselines.

17–24

Ready to orchestrate

Connect agents around one goal, such as bed flow, with a supervisor, gates and a full audit trail.

KEY TAKEAWAYS AT A GLANCE

Forty-five things *to remember*

Three from each of the fifteen chapters, numbered in the order the book makes its case. Tear this out and take it to your next steering committee.

CHAPTER 01 From listening to acting

1. Ambient scribes are infrastructure now: buy on measured time saved, not on demos.
2. The next step is agents that act on the note, with a clinician signing off.
3. Patient trust is the bottleneck: tell patients when AI is listening, and why.

CHAPTER 02 The EHR becomes the platform

4. The EHR is becoming the hub that native, customer-built and outside agents plug into.
5. Judge a system by what can connect to it, with what permissions and audit trail.
6. Let outside AI in read-only first; allow writes only under governance.

CHAPTER 03 The back office pays first

7. Calls, prior authorisation and coding pay back first because their baselines are clean.
8. Coding AI lifts revenue and scrutiny alike: payers are deploying AI of their own.
9. Track abandonment, denials and days to bill before and after go-live.

CHAPTER 04 Hybrid intelligence beats LLM-only

10. Models are becoming a commodity; clinical knowledge, rules and controls are not.
11. Set accuracy bars per field before automating; below them, keep a person in the loop.
12. Prices are falling and open models can run in-country: keep the model layer swappable.

CHAPTER 05 When agents work together

13. Orchestration coordinates agents, people and rules around one hospital goal.
14. Agents are becoming software's main users: booking and data access need agent identity and consent.
15. Use a supervisor with a human escalation path, not a swarm of independent agents.

CHAPTER 06 Beds, theatres and staff

16. Capacity is where orchestration pays first: baselines are clear and results show at the huddle.
17. Prediction helps only when it changes what people do at the morning bed meeting.
18. Backtest on your own data, run in shadow mode, then measure against control wards.

CHAPTER 07 Choosing a hospital system for 2027

19. Your HIS vendor is now your AI vendor: choosing a system is choosing an AI strategy.
20. Built in, bolted on or layered can all work; insist on one record, permission set and audit trail.
21. Score every vendor the same way, with references and certification.

CHAPTER 08 Governance becomes a buying criterion

22. Shadow AI is already here: give staff a sanctioned tool before they reach for public ones.
23. AI earns clinical authority through evidence; accountability stays with people.
24. Ask every vendor for model facts, monitoring and a per-feature off switch.

KEY TAKEAWAYS AT A GLANCE (CONTINUED)

CHAPTER 09 When agents go off-script

25 Agents fail by finding a way: correct answers reached by unauthorised means.

26 Contain agents by design: their own identity, no secrets inside, least privilege, approval for writes.

27 Attackers now move in hours; a 30–60 day patch cycle is a clinical risk.

CHAPTER 10 Capital and consolidation

28 Money is flowing to agents tied to one workflow, and to hospital chains.

29 Expect consolidation: check change-of-control, data-export and audit-log terms before signing.

30 Investors want measurable, profitable outcomes; hospital buyers should demand the same.

CHAPTER 11 India: laying the digital rails

31 India has national rails, but a mid-market that is still largely on paper.

32 DPDP duties phase in through May 2027: build consent and rights handling now.

33 Indian clinical speech is different: ask for local, multilingual evidence.

CHAPTER 12 The GCC: national-scale ambition

34 In the Gulf, governments set the bar through national platforms and AI policy.

35 Abu Dhabi now ties each AI clinical task to evidence and physician oversight.

36 Exchange integration, Arabic and English, and data residency are table stakes.

CHAPTER 13 The India and GCC rulebook

37 In India, the hospital is the data fiduciary: it answers for how its vendors handle patient data.

38 Whether AI is a medical device depends on its intended use, not its technology: get it in writing.

39 In the Gulf, keep health data in-country and insist on live Malaffi, NABIDH or NPHIES links, not roadmaps.

CHAPTER 14 Proving it works

40 Benchmarks rarely test what clinicians use AI for: test on your own encounters.

41 Let systems refuse when unsure, and route uncertain output to people.

42 Measure money and quality together: AI that only lifts coding invites payer pushback.

CHAPTER 15 A 2027 blueprint for hospital leaders

43 Ten requirements define an AI-first hospital system for a 2027 tender.

44 Start where baselines are clean, prove value in 90 days, then scale.

45 Design for the patient's line: AI for logistics, a person accountable for decisions.

KEY TAKEAWAYS AT A GLANCE (CONTINUED)

The next 90 days

DAYS 1 – 30 · LOOK

- List every AI tool in use, approved or not (Ch. 8)
- Give staff a sanctioned assistant (Ch. 8)
- Pick one workflow with a clean baseline (Ch. 3, 6)
- Write the scorecard and record baselines (Ch. 14)

DAYS 31 – 60 · TEST

- Backtest on 12–24 months of your own data (Ch. 14)
- Run in shadow mode, logged but not shown (Ch. 6)
- Set agent identity and approval tiers (Ch. 9)
- Tell patients when AI is listening (Ch. 1)

DAYS 61 – 90 · DECIDE

- Go live in one service, with sign-off (Ch. 5, 6)
- Report gains and side effects monthly (Ch. 14)
- Score vendors on one scorecard (Ch. 7)
- Scale, fix or stop, on the evidence (Ch. 15)

Six questions to take to your board

- Which AI tools are in use today, approved or not, and who owns each one? (Ch. 8)
- Which one number will our first AI project move, and what is its baseline? (Ch. 14)
- What can our AI agents do without a person approving it? (Ch. 9)
- Where is our patient data processed, and who processes it for us? (Ch. 13)
- Are we ready for DPDP duties in May 2027, or for Gulf residency rules? (Ch. 11, 13)
- Would our current hospital system pass the ten requirements for 2027? (Ch. 15)

DATES TO PLAN AROUND

The next 18 months on one line

13 NOV 2026
DPDP consent managers
India[22]

28 DEC 2026
HTI-5 takes effect
US certification rules[38]

1 JAN 2027
Prior-auth APIs due
US payers, CMS-0057-F[120]

MAR 2027
Epic Connectors marketplace
Outside apps in Showroom[3]

13 MAY 2027
DPDP duties and penalties
India[22]

APR 2028
UAE agentic target
Half of government services[16]

IF YOU REMEMBER ONE THING

AI in hospitals is moving from writing things down to doing things. Let it do the work around care, prove every claim on your own data, and keep a named person accountable for every decision it touches.

APPENDIX A

Timeline of key developments

Selected events from 1 August to 4 November 2026, by date of announcement or first report.

August 2026

- **3** Ochsner starts the first Epic Ergo Visit pilot[80]
- **11** Microsoft sets out its agentic AI foundation for healthcare;[93] EMIS announces the name Enlivio Health[220]
- **13** athenahealth opens AI-native encounters to 170,000+ clinicians[19]
- **17** Epic, payers and providers launch real-time prior-auth checks[115]
- **18** Epic rolls out Agent Factory;[87] R1 to acquire Humata Health;[116] Rhapsody payer suite on Epic Showroom[103]
- **19** Oracle adds coding, dictation and chart review to its Clinical AI Agent[74]
- **25** FDB Script Agent's first commercial deployment[76]
- **31** Healthwatch England flags AI scribe errors in NHS records[175]

September 2026

- **2** OLVG results for Epic Art;[9] House VA Committee votes to subpoena Oracle leaders[98]
- **8** OpenAI opens read-only Epic access for ChatGPT for Healthcare;[92] GenHealth.ai raises $16.5M[118]
- **9** VA selects Abridge for enterprise contract;[17] ARPA-H ADVOCATE awards[133]
- **10** WHX Tech publishes $148M GCC healthtech figure[15]
- **14** Oracle agent for nurses;[77] MEDITECH live at 34 HCA hospitals;[18] Kyndryl closes HITL deal;[217] Allina physicians strike[179]
- **15** Imprivata shadow-AI survey;[6] 25 systems at Epic Build-a-Thon[88]
- **16** MRO acquires Vyne Medical[124]
- **17** HISAA cybersecurity bill reintroduced[178]
- **21** DexCare acquires Mila Health[112]
- **22** Heidi raises $340M;[2] Anthropic–OpenEvidence partnership;[131] Epic discloses a development pause[176]
- **23** Oracle plans an Oncology EHR and upstream RCM AI[95]
- **24** BCBSA study on AI-driven coding costs reported[11]
- **25** ABDM reaches 97.8 crore ABHA IDs[7]
- **30** Deadline by which California signs AB 1979 and SB 503[20]

October 2026

- **2** AKASA launches autonomous mid-cycle AI;[121] Texas listing puts Cerner breach at ~20M people[13]
- **3** ChristianaCare goes live on Epic, replacing Oracle Health[107]
- **6** ONC LEAP awards for agentic AI;[145] Ember AI Oracle integration[104]
- **7** Healthleap raises $38M[129]
- **8** Suki Ask launches across four EHRs;[86] UnitedHealthcare's Avery reaches ~21M members[31]
- **12** Northwestern Medicine and OpenEvidence collaborate[102]
- **13** Suki's $90M Series E;[214] Epic leaders discuss protocol agents at Harvard;[79] AMA agentic AI framework;[35] Abridge Assistant[30]
- **14** PHTI: AI scribes and billing tools raise costs[29]
- **16** CERT-In warns of unauthorised AI-agent access[188]
- **19** AHA recommendations to FDA on GenAI devices[189]
- **20** Baba raises $40M;[215] Rhapsody Guardian for agent access;[34] Abu Dhabi Future Health Summit opens[41]
- **21** Epic announces Connectors;[3] Tairex multi-agent system;[134] DoH Abu Dhabi AI Clinician Framework;[36] Tampa General AI-traffic model[143]
- **22** AAIF adds four projects;[33] Microsoft DxO results in *Science*;[24] M42 launches Ossi;[39] Sutter passes 5M AI notes[28]
- **23** Jefferson ambient AI study;[23] Johns Hopkins Bayview AI data notice[26]
- **26** HTI-5 finalised;[21] Global Health Exhibition opens in Riyadh;[239] BILH AI-scribe study reported[27]
- **27** WISeR prior-auth records reported;[171] Oracle pitches AI modules to Epic hospitals[160]
- **28** Zocdoc opens booking to AI agents[25]
- **29** AMA clinical AI liability framework;[190] Pew survey on AI in care[85]
- **30** EY: India healthcare PE/VC $3.1B in Q3[42]
- **31** ORION theatre-scheduling pilot, Manitoba[159]

November 2026

- **2** Hippocratic AI releases Polaris 6 Ultra[136]
- **3** Clinics On Cloud raises ₹120 crore;[216] O-Health funding round[44]
- **4** CERT-In AI Security Auditing Guidelines;[37] Bihar AI call-centre plan reported[32]

APPENDIX B

Glossary

A2A and MCP
Agent2Agent and Model Context Protocol: emerging open protocols for agents to talk to each other and to connect to tools and data.

ABDM
Ayushman Bharat Digital Mission: India's national digital health programme, run by the National Health Authority.

ABHA
Ayushman Bharat Health Account: the health ID that links a person's records across providers under ABDM.

AB-PMJAY
Ayushman Bharat Pradhan Mantri Jan Arogya Yojana: India's publicly funded health insurance scheme.

Agent gateway
A control point between AI agents and hospital systems that checks identity and permissions, removes sensitive data and logs each action.

Agentic AI / AI agent
Software that plans and carries out multi-step tasks across systems, not only answering questions. In this book, agents act within permissions and, for anything that matters, with human sign-off.

Ambient scribe
AI that listens to a clinical conversation and drafts the note for the clinician to review.

Backtest
Running an AI tool on a hospital's past data to see what it would have done, before anyone relies on it.

CDI
Clinical documentation integrity: making sure the record accurately reflects the patient's condition and care.

CDS / CDSS
Clinical decision support (system): tools that give clinicians patient-specific guidance, from rules to AI models.

CERT-In
India's national computer emergency response team. Cyber incidents must be reported to it within six hours.

Command centre
A hub where staff and software coordinate beds, transfers and flow across a hospital or group.

Consent manager
Under India's DPDP Act, a registered platform through which people give, manage and withdraw consent across organisations.

Data fiduciary
Under the DPDP Act, the organisation that decides why and how personal data is processed. For patient data, the hospital; vendors are its processors.

DPDP
India's Digital Personal Data Protection Act, 2023, and its Rules, 2025.

ED boarding
Admitted patients waiting in the emergency department because no inpatient bed is free.

EHR / EMR
Electronic health record / electronic medical record. Often used interchangeably; an EHR typically spans providers.

FHIR
Fast Healthcare Interoperability Resources: the HL7 standard for exchanging health data through APIs.

HIS
Hospital information system: the software that runs a hospital's clinical, administrative and financial work, usually including the EMR.

HTI-5
A US federal health IT rule from ASTP/ONC, finalised in October 2026.

Malaffi and NABIDH
The health information exchanges of Abu Dhabi and Dubai.

NABH
National Accreditation Board for Hospitals & Healthcare Providers (India).

NHA / NHCX
India's National Health Authority, and its National Health Claims Exchange for insurance claims.

NPHIES
Saudi Arabia's national platform for health information exchange and insurance transactions.

APPENDIX B · GLOSSARY (CONTINUED)

Objective element
The smallest assessable requirement in NABH standards; the 6th edition has 639.

Open-weight model
An AI model whose parameters are published, so it can be run on a hospital's or partner's own servers.

Orchestration
Coordinating several AI agents, people and systems towards one goal, with a supervisor, rules and an audit trail.

PDPL
Saudi Arabia's Personal Data Protection Law, which limits transfers of personal data abroad.

Prior authorisation (pre-auth)
Approval an insurer or TPA must give before certain treatments are covered.

Prompt injection
Instructions hidden in a document, email or web page that make an AI system do something its user did not intend.

RCM
Revenue cycle management: registration, eligibility, coding, billing, claims and collections.

SaMD
Software as a medical device: software intended for diagnosis, treatment or related clinical purposes, regulated as a device (CDSCO in India, SFDA in Saudi Arabia).

Shadow AI
AI tools used in an organisation without formal approval from IT or governance.

Shadow mode
Running an AI tool in the background so its predictions can be checked before anyone acts on them.

Significant Data Fiduciary
An organisation designated by the Indian government for extra DPDP duties: a data protection officer, impact assessments and annual data audits.

SNOMED CT, LOINC, ICD-10
Coding systems for clinical terms, laboratory tests and diagnoses.

TEFCA
Trusted Exchange Framework and Common Agreement: the US national framework for health data exchange.

TPA
Third-party administrator: processes health insurance claims on behalf of insurers, common in India.

WER
Word error rate: the share of words a speech-recognition system gets wrong.

Zero data retention
An AI provider setting under which request inputs and outputs are not stored after processing or used for training.

APPENDIX C

References

Sources are numbered in order of first citation. Dates are publication dates as shown by the source; undated pages give the date we accessed them. All links were checked between 2 and 9 November 2026.

1 **Surgy Health (S. Nagekar).** "Hospital AI Budget 2027: Where to Invest and How to Choose." 28 Oct 2026.
https://surgy.health/blog/ai-resource-allocation-hospitals-2027

2 **HIT Consultant.** "Heidi Secures $340M to Transition from Clinical Documentation to Supervised Agentic Execution." 22 Sep 2026.
https://hitconsultant.net/2026/09/22/heidi-secures-340m-series-c-general-catalyst-ai-care-partner-clinical-agents/

3 **Becker's Hospital Review.** "Epic takes aim at 'swivel-chair' work with AI Connectors." 21 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/epic-takes-aim-at-swivel-chair-work-with-ai-connectors/

4 **Becker's Hospital Review.** "12 health systems expanding their use of agentic AI." 31 Jul 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ai/12-health-systems-expanding-their-use-of-agentic-ai/

5 **International Business Times UK.** "Palantir AI Software Saved Nearly 900 Lives at One Florida Hospital by Catching Sepsis Before Doctors Could." 10 Jun 2026.
https://www.ibtimes.co.uk/ai-tool-saves-lives-detecting-sepsis-florida-hospital-1801840

6 **HIT Consultant.** "Imprivata Report Reveals 72% of Healthcare Organizations Run Unapproved AI as Autonomous Agents Enter Clinical Care." 15 Sep 2026.
https://hitconsultant.net/2026/09/15/imprivata-report-healthcare-agentic-ai-trust-gap-identity-led-governance-shadow-ai/

7 **DD India (Doordarshan).** "Arogya Manthan 2026: J.P. Nadda highlights expansion of AB PM-JAY and digital health ecosystem." 25 Sep 2026.
https://ddindia.co.in/2026/09/arogya-manthan-2026-j-p-nadda-highlights-expansion-of-pm-jay-and-digital-health-ecosystem/

8 **Nextgov/FCW.** "VA to gradually expand medical AI scribe following successful pilots." 14 Sep 2026.
https://www.nextgov.com/artificial-intelligence/2026/09/va-gradually-expand-medical-ai-scribe-following-successful-pilots/415971/

9 **Epic.** "AI Helps OLVG Clinicians Reduce Note-Writing Time By Over 30%." 2 Sep 2026.
https://www.epic.com/epic/post/ai-helps-olvg-clinicians-reduce-note-writing-time-by-over-30/

10 **Becker's Hospital Review.** "Epic's market share over the last 5 years." 14 May 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/epics-market-share-over-the-last-5-years/

11 **Becker's Hospital Review.** "Hospitals say they're losing the AI billing war. A new BCBS study suggests otherwise." 24 Sep 2026.
https://beckershospitalreview.com/payer/hospitals-say-theyre-losing-the-ai-billing-war-a-new-bcbs-study-suggests-otherwise

12 **HIT Consultant.** "Capstone Partners Reports 93.8% Surge in AI-Enabled Healthcare IT M&A." 29 Sep 2026.
https://hitconsultant.net/2026/09/29/capstone-partners-report-ai-enabled-healthcare-it-ma/

13 **eSecurity Planet.** "Texas AG Says Oracle Health's 2025 Breach Affected 20M Individuals." 7 Oct 2026.
https://esecurityplanet.com/news/oracle-health-breach-20-million-texas-ag

14 **Indian Pharma Post (IPP Bureau).** "AI adoption in Indian hospitals remains largely at pilot stage: Bain-HealthQuad report." 9 Sep 2026.
https://www.indianpharmapost.com/hospitals/ai-adoption-in-indian-hospitals-remains-largely-at-pilot-stage-bain-healthquad-report-21769

15 **WHX Tech / World Health Expo (Informa Markets) press release.** "US$148m healthtech investment puts GCC startups in focus at WHX Tech." 10 Sep 2026.
https://www.worldhealthexpo.com/events/healthcare/tech/resources/press-media-centre/press-releases/us148m-healthtech-investment-puts-gcc-startups-in-focus-at-whx-tech/

16 **Dubai Media Office.** "Mohammed bin Rashid unveils framework to deploy agentic AI across 50% of government sectors within 2 years." 23 Apr 2026.
https://mediaoffice.ae/en/news/2026/april/23-04/mohammed-bin-rashid-chairs-uae-cabinet-meeting

17 **Abridge press release (Business Wire).** "Abridge Selected to Bring Ambient AI to Veteran Care Nationwide Through the VA Enterprise Contract." 9 Sep 2026.
https://www.businesswire.com/news/home/20260909133398/en/Abridge-Selected-to-Bring-Ambient-AI-to-Veteran-Care-Nationwide-Through-the-VA-Enterprise-Contract

18 **HIT Consultant.** "MEDITECH Completes Expanse Implementation Across 34 Hospitals in HCA Healthcare's First Two U.S. Divisions." 16 Sep 2026.
https://hitconsultant.net/2026/09/16/meditech-completes-expanse-ehr-implementation-34-hospitals-hca-healthcare-divisions/

19 **athenahealth (Business Wire, via 01net.it).** "athenahealth's AI-native athenaOne® experiences now available to 170,000 clinicians." 13 Aug 2026.
https://www.01net.it/athenahealths-ai-native-athenaone-experiences-now-available-to-170000-clinicians/

20 **California Health Care Foundation (CHCF).** "Newsom Approves 3 Health Care AI Bills, Vetoes 2." 8 Oct 2026.
https://chcf.org/blog/2026/10/08/newsom-approves-3-health-care-ai-bills-vetoes-2

21 **ASTP/ONC (HealthIT.gov).** "HTI-5 Final Rule." 26 Oct 2026.
https://www.healthit.gov/hti5

22 **Bar & Bench (S.S. Rana & Co. viewpoint).** "MeitY notifies final Digital Personal Data Protection Rules, 2025." 17 Nov 2025.
https://www.barandbench.com/amp/story/view-point/meity-notifies-final-digital-personal-data-protection-rules-2025

23 **Becker's Hospital Review.** "Jefferson Health's ambient AI users closed charts up to 39% faster." 23 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ai/jefferson-healths-ambient-ai-users-closed-charts-up-to-39-faster/

24 **Becker's Hospital Review.** "Microsoft AI diagnostic tool tops generalist physicians, 83% to 36%." 22 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/digital-health/microsoft-ai-diagnostic-tool-tops-generalist-physicians-83-to-36/

25 **Healthcare Dive.** "Zocdoc invites AI agents to book appointments on behalf of patients." 28 Oct 2026.
https://healthcaredive.com/news/zocdoc-invites-ai-agents-to-book-appointments-on-behalf-of-patients/873035

26 **WMAR-2 News.** "Johns Hopkins alerts 3,000+ patients after data sent to AI tool." 23 Oct 2026.
https://www.wmar2news.com/local/patients-data-sent-to-ai-tool-johns-hopkins-bayview-medical-center-oct-23-2026

27 **Becker's Hospital Review.** "Heavy AI scribe users at Beth Israel Lahey cut note time 2.5 minutes." 26 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ai/heavy-ai-scribe-users-at-beth-israel-lahey-cut-note-time-2-5-minutes/

28 **Becker's Hospital Review.** "Sutter Health reports 5 million AI-generated clinical notes." 22 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ai/sutter-health-reports-5-million-ai-generated-clinical-notes/

29 **Becker's Hospital Review.** "AI scribes, billing tools drive up healthcare costs: Report." 14 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/digital-health/ai-scribes-billing-tools-drive-up-healthcare-costs-report/

30 **Abridge.** "Meet the AI Assistant for Healthcare." 13 Oct 2026.
https://www.abridge.com/blog/abridge-assistant-is-the-ai-assistant-for-healthcare

APPENDIX C · REFERENCES (CONTINUED)

31 **Becker's Payer Issues.** "UnitedHealthcare's AI assistant expands to 21 million members." 8 Oct 2026.
https://www.beckerspayer.com/virtual-care/unitedhealthcares-ai-assistant-expands-to-21-million-members/

32 **Digital Health News.** "Bihar Plans AI Call Centre for Government Hospital OPD Slots." 4 Nov 2026.
https://www.digitalhealthnews.com/bihar-plans-ai-call-centre-to-help-patients-book-govt-hospital-opd-slots

33 **SiliconANGLE.** "Agentic AI Foundation adds four more open-source AI tools." 22 Oct 2026.
https://siliconangle.com/2026/10/22/agentic-ai-foundation-adds-four-more-open-source-ai-tools/

34 **PR Newswire (Rhapsody).** "Rhapsody Launches Guardian to Govern AI Agent and API Access to Healthcare Data and Services." 20 Oct 2026.
https://www.prnewswire.com/news-releases/rhapsody-launches-guardian-to-govern-ai-agent-and-api-access-to-healthcare-data-and-services-302913231.html

35 **Becker's Hospital Review.** "AMA: Agentic AI must earn clinical authority through evidence." 13 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/digital-health/ama-agentic-ai-must-earn-clinical-authority-through-evidence/

36 **TradeArabia.** "DoH publishes AI Clinician Framework." 21 Oct 2026.
https://tradearabia.com/News/488291/DoH-publishes-AI-Clinician-Framework

37 **Cloud Security Alliance (CSA Labs research note).** "Auditable Agents: CERT-In's AI Auditing Guidelines and Global Assurance." 6 Nov 2026.
https://labs.cloudsecurityalliance.org/research/csa-research-note-certin-ai-security-auditing-guidelines-202/

38 **Medical Daily.** "Hospitals Lose Federal AI Model Card Disclosures Under HTI-5 as Epic Pledges to Keep Its Own Transparency Tools." 3 Nov 2026.
https://www.medicaldaily.com/federal-ai-transparency-rules-health-records-hti-5-480729

39 **The National.** "Abu Dhabi's M42 launches AI platform to give doctors more time with patients." 22 Oct 2026.
https://www.thenationalnews.com/news/uae/2026/10/22/abu-dhabi-m42-ossi-ai-doctors/

40 **Gulf News.** "Abu Dhabi hospitals to use AI to take paperwork out of consultation room." 23 Oct 2026.
https://gulfnews.com/uae/abu-dhabi/abu-dhabi-hospitals-to-use-ai-to-take-paperwork-out-of-consultation-room-1.500720169

41 **Gulf Daily News.** "Future Health to embed wearables into preventive healthcare." 21 Oct 2026.
https://www.gdnonline.com/Details/1410374/Future-Health-to-embed-wearables-into-preventive-healthcare

42 **EY India.** "PE/VC investment activity in India reached US$17.1 billion across 348 deals in 3Q2026, a 46% year-on-year increase." 30 Oct 2026.
https://www.ey.com/en_in/newsroom/2026/10/pe-vc-investment-activity-in-india-reached-us-dollor-17-point-1-billion-across-348-deals-in-3q2026-a-46-percent-year-on-year-increase

43 **arXiv.** "Evaluating Ambient Clinical Scribes in India: The Need for Multilingual Real-World Clinical Conversation Data." 15 Sep 2026.
https://arxiv.org/pdf/2609.17355

44 **Dealroom News.** "Apollo's Sangita Reddy backs O-Health's edge AI for doctors." 3 Nov 2026.
https://dealroom.co/news/169033-apollos-sangita-reddy-backs-o-healths-edge-ai-for-doctors/

45 **Poon et al., Journal of the American Medical Informatics Association.** 2025.
https://academic.oup.com/jamia/article/32/7/1093/8125015

46 **American Medical Association.** "More than 80% of physicians use AI professionally: AMA survey." 2026.
https://ama-assn.org/practice-management/digital-health/more-80-physicians-use-ai-professionally-ama-survey

47 **American Medical Association.** "2 in 3 physicians are using health AI, up 78% from 2023." 2025.
https://www.ama-assn.org/practice-management/digital-health/2-3-physicians-are-using-health-ai-78-2023

48 **Surgy Health.** "AI in Healthcare Statistics 2026: 100+ Data Points from Primary Sources." 27 Oct 2026.
https://surgy.health/blog/ai-in-healthcare-statistics

49 **ASTP/ONC.** "Hospital trends in the use, evaluation and governance of predictive AI, 2023–2024 (Data Brief No. 80)." 2025.
https://healthit.gov/data/data-briefs/hospital-trends-use-evaluation-and-governance-predictive-ai-2023-2024/

50 **Everson et al., JAMA Network Open.** 2025.
https://pmc.ncbi.nlm.nih.gov/articles/PMC12701511

51 **Elsevier.** "Global study of clinicians finds nurses being left out of clinical AI adoption (Clinician of the Future 2026)." 2026.
https://www.elsevier.com/en-gb/about/press-releases/global-study-of-clinicians-by-elsevier-finds-nurses-being-left-out-of-clinical-ai-adoption

52 **Hernström et al., The Lancet Digital Health.** 2025.
https://doi.org/10.1016/S2589-7500%2824%2900267-X

53 **Adams et al., Nature Medicine.** 2022.
https://doi.org/10.1038/s41591-022-01894-0

54 **Lin et al., Nature Medicine.** 2024.
https://doi.org/10.1038/s41591-024-02961-4

55 **Martinez-Gutierrez et al., JAMA Neurology.** 2023.
https://doi.org/10.1001/jamaneurol.2023.3206

56 **Olson et al., JAMA Network Open.** 2025.
https://doi.org/10.1001/jamanetworkopen.2025.34976

57 **Wong et al., JAMA Internal Medicine.** 2021.
https://doi.org/10.1001/jamainternmed.2021.2626

58 **Rotenstein et al., JAMA.** 2026.
https://doi.org/10.1001/jama.2026.2253

59 **Asgari et al., npj Digital Medicine.** 2025.
https://doi.org/10.1038/s41746-025-01670-7

60 **Goh et al., JAMA Network Open.** 2024.
https://doi.org/10.1001/jamanetworkopen.2024.40969

61 **Bedi et al., JAMA.** 2025.
https://pmc.ncbi.nlm.nih.gov/articles/PMC11480901

62 **US Food and Drug Administration.** "Artificial Intelligence-Enabled Medical Devices (list, 4 September 2026 update)." 4 Sep 2026.
https://www.fda.gov/medical-devices/artificial-intelligence-enabled-medical-devices/list-artificial-intelligence-enabled-medical-devices

63 **Windecker et al., JAMA Network Open.** 2025.
https://doi.org/10.1001/jamanetworkopen.2025.8052

64 **MarketsandMarkets.** "Artificial Intelligence (AI) in Healthcare market report." 2025.
https://www.marketresearch.com/MarketsandMarkets-v3719/Artificial-Intelligence-AI-Healthcare-Offering-41091302/

65 **Fortune Business Insights.** "Artificial Intelligence in Healthcare Market." 2026.
https://www.fortunebusinessinsights.com/industry-reports/artificial-intelligence-in-healthcare-market-100534

66 **Menlo Ventures.** "2025: The State of AI in Healthcare." 2025.
https://menlovc.com/perspective/2025-the-state-of-ai-in-healthcare/

67 **Rock Health.** "2025 year-end digital health funding overview: a tale of two markets." 2026.
https://rockhealth.com/insights/2025-year-end-digital-health-funding-overview-a-tale-of-two-markets/

APPENDIX C · REFERENCES (CONTINUED)

68 **Bain & Company.** "The 2026 Healthcare AI ROI Scorecard." 2026.
https://www.bain.com/insights/the-2026-healthcare-ai-roi-scorecard/

69 **Nong and Platt, JAMA Network Open.** 2025.
https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2830240

70 **Philips.** "Building trust in healthcare AI: five key insights from the 2025 Future Health Index." 2025.
https://www.philips.com/a-w/about/news/archive/blogs/innovation-matters/2025/building-trust-in-healthcare-ai-five-key-insights-from-the-2025-future-health-index

71 **World Health Organization.** "One in four doctors nearing retirement age as WHO warns of global health workforce gap." 18 Sep 2026.
https://www.who.int/news/item/18-09-2026-one-in-four-doctors-nearing-retirement-age-as-who-warns-of-global-health-workforce-gap

72 **American Medical Association.** "AMA survey: prior authorization reform pledge falls short for physicians." 2026.
https://www.ama-assn.org/press-center/ama-press-releases/ama-survey-prior-authorization-reform-pledge-falls-short-physicians

73 **HIT Consultant.** "Heidi II launches autonomous clinical agents to execute 10M weekly visit tasks." 2 Oct 2026.
https://hitconsultant.net/2026/10/02/heidi-ii-launches-autonomous-clinical-agents-to-execute-10m-weekly-visit-tasks/

74 **Oracle.** "Oracle Health Expands Clinical AI Agent with Coding, Dictation, and Chart Review." 19 Aug 2026.
https://www.oracle.com/news/announcement/oracle-health-expands-clinical-ai-agent-with-automated-coding-dictation-chart-review-2026-08-19/

75 **Healthcare IT News.** "Private Indian hospital takes platform-first approach to ambient AI." Undated; accessed Nov 2026.
https://www.healthcareitnews.com/news/asia/private-indian-hospital-takes-platform-first-approach-ambient-ai

76 **FDB/Tebra press release (PR Newswire, via TechEdge AI).** "FDB and Tebra Bring AI From Clinical Notes Into Prescription Workflows." 25 Aug 2026.
https://techedgeai.com/fdb-and-tebra-bring-ai-from-clinical-notes-into-prescription-workflows/

77 **Oracle.** "Oracle Health Clinical AI Agent Helps Nurses Alleviate Documentation Burden and Streamline Care." 14 Sep 2026.
https://www.oracle.com/news/announcement/oracle-health-clinical-ai-agent-helps-nurses-alleviate-documentation-burden-and-streamline-care-2026-09-14/

78 **Becker's Hospital Review.** "HCA expands AI nurse handoff tool to 15 more hospitals." 29 Sep 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/innovation/hca-expands-ai-nurse-handoff-tool-to-15-more-hospitals/

79 **Becker's Hospital Review.** "Epic explores AI agents for routine visits." 20 Oct 2026.
https://www.beckershospitalreview.com/uncategorized/epic-explores-ai-agents-for-routine-visits/

80 **Healthcare Dive.** "Epic targets outpatient visits with new AI tool." 19 Aug 2026.
https://healthcaredive.com/news/epic-targets-outpatient-visits-with-new-ai-tool/828230

81 **IKS Health.** "IKS Health Implements Scribble Select, an Ambient Scribe Interface, at Western Washington Medical Group." 19 Aug 2026.
https://ikshealth.com/?p=28537

82 **Tripura Star News.** "Baby Memorial Hospital Group partners with DeepInfinity.AI to scale healthcare AI." 30 Oct 2026.
https://www.tripurastarnews.com/baby-memorial-hospital-group-partners-with-deepinfinity-ai-to-scale-healthcare-ai/

83 **Stacker.** "The doctor's AI will see you now." 6 Oct 2026.
https://kioncentralcoast.com/news/2026/10/06/the-doctors-ai-will-see-you-now/

84 **The Cavalier Daily.** "U.Va. Health cuts scribe hours as AI documentation tools expand across health system." 1 Oct 2026.
https://cavalierdaily.com/article/2026/10/uva-health-cuts-scribe-hours-as-ai-documentation-tools-expand-across-health-system

85 **Pew Research Center.** "Americans are open to AI in health care but want humans making final calls." 29 Oct 2026.
https://www.pewresearch.org/science/2026/10/29/americans-are-open-to-ai-in-healthcare-but-want-humans-making-final-calls.md

86 **Suki (press release, via Barchart/Webull).** "Suki Launches "Ask," Accelerating Complex Clinical Decision Support Powered by Wolters Kluwer UpToDate Expert AI." 8 Oct 2026.
https://www.webull.com/news/15712890866582528

87 **Fierce Healthcare.** "Epic expands AI ambitions with agent platform, Cosmos-powered predictions and workflow automation." Aug 2026.
https://fiercehealthcare.com/health-tech/epic-expands-ai-ambitions-agent-platform-cosmos-powered-predictions-and-deeper-workflow

88 **Becker's Hospital Review.** "25 health systems build AI agents with Epic." 15 Sep 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/25-health-systems-build-ai-agents-with-epic/

89 **Epic.** "Introducing Connectors: A New Way for Apps to Communicate With Epic." 21 Oct 2026.
https://www.epic.com/epic/post/introducing-connectors-a-new-way-for-apps-to-communicate-with-epic/

90 **Epic.** "Emmie (product page)." Undated; accessed Nov 2026.
https://www.epic.com/software/emmie

91 **Becker's Hospital Review.** "The EHR is becoming a front door for healthcare AI." 22 Sep 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/the-ehr-is-becoming-healthcare-ais-front-door/

92 **OpenAI.** "Healthcare organizations can now connect EHR and additional industry data to ChatGPT." 8 Sep 2026.
https://openai.com/index/chatgpt-connects-health-records-and-healthcare-sources/

93 **Microsoft (Microsoft Cloud Blog).** "Building the foundation for agentic AI in healthcare." 11 Aug 2026.
https://www.microsoft.com/en-us/microsoft-cloud/blog/healthcare/2026/08/11/building-the-foundation-for-agentic-ai-in-healthcare/

94 **Becker's Hospital Review.** "Franciscan, Brown Health scale Microsoft's agentic AI tools." 12 Aug 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/innovation/franciscan-brown-health-scale-microsofts-agentic-ai-tools/

95 **Oracle.** "Oracle to Advance Cancer Care with New Oncology EHR." 23 Sep 2026.
https://www.oracle.com/news/announcement/oracle-to-advance-cancer-care-with-new-oncology-ehr-2026-09-23/

96 **HIT Consultant.** "Oracle Health unveils upstream AI revenue cycle management portfolio." 23 Sep 2026.
https://hitconsultant.net/2026/09/23/oracle-health-unveils-upstream-ai-revenue-cycle-management-portfolio/

97 **Oracle.** "Oracle PartnerNetwork Expands Incentives to Help Partners Accelerate Healthcare Transformation." 22 Sep 2026.
https://www.oracle.com/news/announcements/oracle-partner-network-expands-incentives-to-help-partners-accelerate-healthcare-transformation-2026-09-22/

98 **FedScoop.** "House committee votes to subpoena Oracle after fiery hearing on VA's $27B EHRM contract." 2 Sep 2026.
https://fedscoop.com/oracle-subpoena-house-hearing-va-27b-ehrm-contract/

99 **MEDITECH.** "MEDITECH recognized as a Leader in the 2026 Gartner Magic Quadrant for Enterprise Electronic Health Records." 25 Aug 2026.
https://ehr.meditech.com/news/meditech-recognized-as-a-leader-in-the-2026-gartner-magic-quadrant-for-enterprise-electronic

100 **Avo (press release, syndicated via cnhinews.com).** "Avo Expands MEDITECH Alliance Partnership, Bringing Its Suite of AI-Powered Products to MEDITECH's EHR." 9 Sep 2026.
https://www.cnhinews.com/news/article_b9843f34-4c73-5258-9f61-366a60238914.html

101 **Memorial Sloan Kettering Cancer Center.** "MSK and OpenEvidence partner to advance precision oncology at point of care." 16 Sep 2026.
https://www.mskcc.org/news-releases/mskcc-msk-and-openevidence-partner-to-advance-precision-oncology-at-point-of-care

102 **OpenEvidence/Northwestern Medicine press release (Business Wire, via Webull).** "OpenEvidence and Northwestern Medicine Launch Collaboration to Integrate Clinical Evidence into the EHR." 12 Oct 2026.
https://www.webull.com/news/15736327729259520

103 **Rhapsody.** "Rhapsody Brings AI-Powered Payer Interoperability Suite to Connection Hub on Epic Showroom, Accelerating Integration for Epic Payer Customers." 18 Aug 2026.
https://rhapsody.health/blog/rhapsody-brings-ai-powered-payer-interoperability-suite-to-connection-hub-on-epic-showroom-accelerating-integration-for-epic-payer-customers/

104 **Ember AI press release (PR Newswire).** "Ember AI Launches Strategic Partner Program and New Oracle Health Integration." 6 Oct 2026.
https://www.prnewswire.com/news-releases/ember-ai-launches-strategic-partner-program-and-new-oracle-health-integration-302899844.html

APPENDIX C · REFERENCES (CONTINUED)

105 **Aiva Health (PR Newswire, via yourwyominglink.com).** "Aiva's turn-key MCP Connector and A2A Agent enables thousands of clinical use cases." 18 Aug 2026.
https://www.yourwyominglink.com/aiva-leverages-chatgpt-for-healthcare-claude-for-healthcare-and-gemini-enterprise-to-control-40-hospital/article_39501d87-5700-5282-82ad-939fc1960dfe.html

106 **Healthcare IT News.** "AI should make healthcare more humane, Taiwanese health leader says." Aug 2026.
https://www.healthcareitnews.com/news/asia/ai-should-make-healthcare-more-humane-taiwanese-health-leader-says

107 **Becker's Hospital Review.** "2026's biggest EHR switches, by vendor." 7 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/2026s-biggest-ehr-switches-by-vendor/

108 **Becker's Hospital Review.** "Northwell completes Epic scheduling conversion for 4 hospitals, 538 sites." 21 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/northwell-completes-epic-scheduling-conversion-for-4-hospitals-538-sites/

109 **Modern Healthcare.** "How AI could make it harder for hospitals to switch EHR vendors." Undated; accessed Nov 2026.
https://www.modernhealthcare.com/health-tech/mh-ai-hospitals-ehr-epic-oracle/

110 **MobiHealthNews.** "Tampa General Hospital taps Hyro voice AI to boost call center efficiency." Undated; accessed Nov 2026.
https://www.mobihealthnews.com/news/tampa-general-hospital-taps-hyro-voice-ai-boost-call-center-efficiency

111 **WellSpan Health (GlobeNewswire).** "WellSpan expands Hippocratic AI partnership to support clinical operations and improve patient experience." 30 Jul 2026.
https://www.globenewswire.com/news-release/2026/07/30/3336096/0/en/WellSpan-expands-Hippocratic-AI-partnership-to-support-clinical-operations-and-improve-patient-experience.htm

112 **HIT Consultant.** "DexCare Acquires AI Care Coordination Platform Mila Health to Unify Patient Access and Conversational Scheduling." 21 Sep 2026.
https://hitconsultant.net/2026/09/21/dexcare-acquires-mila-health-ai-care-coordination-autonomous-patient-access/

113 **Nextech press release (GlobeNewswire, via TechDogs).** "Nextech Unveils Agentic AI Capabilities for Patient Experience through Partnerships with Assort Health and EliseAI." 9 Oct 2026.
https://www.techdogs.com/tech-news/globe-newswire/nextech-unveils-agentic-ai-capabilities-for-patient-experience-through-partnerships-with-assort-health-and-eliseai

114 **Assort Health press release.** "Assort Health Raises $120 Million Series C to Scale Largest Deployment of AI Agents for the Patient Journey." 24 Jun 2026.
https://aijourn.com/assort-health-raises-120-million-series-c-to-scale-largest-deployment-of-ai-agents-for-the-patient-journey/

115 **Epic.** "Payers, Providers, and Epic Launch Real-Time Prior Authorization Checks." 17 Aug 2026.
https://www.epic.com/epic/post/payers-providers-and-epic-launch-real-time-prior-authorization-checks/

116 **TechTarget (RevCycle Management).** "R1 to acquire Humata Health for AI-driven prior authorizations." 18 Aug 2026.
https://www.techtarget.com/revcyclemanagement/news/366648923/R1-to-acquire-Humata-Health-for-AI-driven-prior-authorizations

117 **HIT Consultant.** "R1 Unveils Phare Utilization Management on Phare OS to Eliminate Medical Necessity Denials." 16 Sep 2026.
https://hitconsultant.net/2026/09/16/r1-announces-phare-utilization-management-phare-os-medical-necessity-denials-guthrie-clinic/

118 **GenHealth.ai press release (Newswire).** "GenHealth.ai Raises $16.5 Million Series A to Build AI Agents That Run the Medical Back Office." 8 Sep 2026.
https://www.newswire.com/news/genhealth-ai-raises-16-5-million-series-a-to-build-ai-agents-that-run-the

119 **Bloomberg Government.** "Matt Holt's Thoreau Invests in Tech Platform Penelope Health." 18 Sep 2026.
https://news.bgov.com/private-equity/matt-holts-thoreau-invests-in-tech-platform-penelope-health

120 **CMS.** "CMS Interoperability and Prior Authorization Final Rule CMS-0057-F (fact sheet)." 17 Jan 2024.
https://www.cms.gov/newsroom/fact-sheets/cms-interoperability-prior-authorization-final-rule-cms-0057-f

121 **HIT Consultant.** "AKASA Launches Autonomous Mid-Cycle AI for Inpatient Coding and CDI." 2 Oct 2026.
https://hitconsultant.net/2026/10/02/akasa-launches-autonomous-ai-mid-cycle-platform-inpatient-coding-cdi/

122 **MEDITECH.** "MEDITECH's latest AI innovations support smarter care and stronger performance." 10 Mar 2026.
https://ehr.meditech.com/news/meditech-s-latest-ai-innovations-support-smarter-care-and-stronger-performance

123 **WellSky (Business Wire).** "WellSky® Unveils Agentic AI Workflows to Help Skilled Nursing Providers Turn Data Into Action Faster." 22 Sep 2026.
https://www.businesswire.com/news/home/20260922162212/en/WellSky-Unveils-Agentic-AI-Workflows-to-Help-Skilled-Nursing-Providers-Turn-Data-Into-Action-Faster

124 **HIT Consultant.** "M&A: MRO Acquires Vyne Medical to Expand Intelligent Healthcare Intake Automation." 16 Sep 2026.
https://hitconsultant.net/2026/09/16/mro-acquires-vyne-medical-intelligent-document-processing-clinical-data-management/

125 **TechCrunch.** "Insurers claim AI is already increasing healthcare costs." 26 Sep 2026.
https://techcrunch.com/2026/09/26/insurers-claim-ai-is-already-increasing-healthcare-costs/

126 **MedCity News.** "Parakeet Health secures $10M for conversational agentic platform." 2 Oct 2026.
https://medcitynews.com/2026/10/parakeet-health-secures-10m-for-conversational-agentic-platform/

127 **Healthcare Dive.** "Zocdoc acquires Sohar Health in latest push to scale tech infrastructure." 15 Oct 2026.
https://www.healthcaredive.com/news/zocdoc-acquires-sohar-health-in-latest-push-to-scale-tech-infrastructure/833129/

128 **Becker's Hospital Review.** "The pilot we killed: Why Nebraska Medicine shelved its AI cancer registry tool." 21 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ai/the-pilot-we-killed-why-nebraska-medicine-shelved-its-ai-cancer-registry-tool/

129 **TechCrunch.** "Healthleap raises $38M for its AI that flags hospital patients who may need a closer look." 7 Oct 2026.
https://techcrunch.com/?p=3174992

130 **Healthcare.Digital (Nelson Advisors).** "OpenEvidence secures $250 Million alternative financing round led by Andreessen Horowitz and Byers Capital at $15 Billion valuation." 25 Sep 2026.
https://www.healthcare.digital/single-post/nelson-advisors-openevidence-secures-250-million-alternative-financing-round-led-by-andreessen-hor

131 **Reuters.** "Exclusive-Anthropic, OpenEvidence partner to bring medical AI worldwide." 23 Sep 2026.
https://www.thestar.com.my/tech/tech-news/2026/09/23/exclusive-anthropic-openevidence-partner-to-bring-medical-ai-worldwide

132 **Digital Health News.** "Govt to Deploy AI-Powered Clinical Decision Support System Across 70,000 Hospitals." 24 Dec 2025.
https://www.digitalhealthnews.com/govt-to-deploy-ai-powered-clinical-decision-support-system-across-70-000-hospitals

133 **ARPA-H.** "ARPA-H launches the world's first bid to build an FDA-authorized clinical AI for cardiovascular care." 9 Sep 2026.
https://arpa-h.gov/news-and-events/arpa-h-launches-worlds-first-bid-build-fda-authorized-clinical-ai-cardiovascular

134 **Tairex press release (PR Newswire Asia, via Kyodo News PR Wire).** "Tairex Launches Multi-Agent AI System to Transform Clinical Collaboration and Decision Support." 22 Oct 2026.
https://kyodonewsprwire.jp/release/202610228537

135 **Healthcare IT Today.** "MRO's Prodigy brings AI-human collaboration to clinical data abstraction." 21 Oct 2026.
https://www.healthcareittoday.com/2026/10/21/mros-prodigy-brings-ai-human-collaboration-to-clinical-data-abstraction/

136 **HIT Consultant.** "Hippocratic AI launches Polaris 6 Ultra." 2 Nov 2026.
https://hitconsultant.net/?p=98478

137 **Anthropic.** "Claude Haiku 5.5." 7 Oct 2026.
https://www.anthropic.com/claude-haiku-5-5

138 **Unite.AI.** "OpenAI unveils GPT-6.1 Sol at DevDay." 29 Sep 2026.
https://www.unite.ai/openai-unveils-gpt-6-1-sol-at-devday-with-new-codex-and-chatgpt-tools/

139 **Anthropic.** "Claude Sonnet 5.5." 28 Sep 2026.
https://www.anthropic.com/claude-sonnet-5-5

140 **Let's Data Science.** "Mistral releases open weights for 1-trillion-parameter Mistral Large 4." 29 Oct 2026.
https://letsdatascience.com/news/mistral-releases-open-weights-for-1-trillion-parameter-mistr-d712eb0e

141 **SiliconANGLE.** "Moonshot AI launches Kimi K3.1 Pro." 23 Oct 2026.
https://siliconangle.com/2026/10/23/moonshot-ai-launches-new-flagship-kim-k3-1-pro-model-with-3t-parameters/

APPENDIX C · REFERENCES (CONTINUED)

142 **How2Shout.** "Nvidia's new speech model cut errors on Indian doctors' dictation by 62%." 28 Oct 2026.
https://www.how2shout.com/news/nvidias-new-speech-model-cut-errors-on-indian-doctors-dictation-by-62.html

143 **Becker's Hospital Review.** "Tampa General trains Nvidia model to police its own AI traffic." 21 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ai/tampa-general-trains-nvidia-model-to-police-its-own-ai-traffic/

144 **TechCrunch.** "Anthropic officially files for its public debut." 26 Oct 2026.
https://techcrunch.com/2026/10/26/anthropic-officially-files-for-its-public-debut/

145 **ASTP/ONC (HealthIT.gov).** "HHS Announces 2026 LEAP in Health IT Awardees to Accelerate the Use of Agentic AI in Clinical Care and Close Gaps in Laboratory Interoperability." 6 Oct 2026.
https://healthit.gov/?p=217574

146 **BleepingComputer.** "Nearly 700 rogue AI agents coordinated in the Hugging Face attack." 27 Aug 2026.
https://bleepingcomputer.com/news/security/nearly-700-rogue-ai-agents-coordinated-in-the-hugging-face-attack

147 **OpenAI.** "Hugging Face incident: technical report." Jul 2026.
https://cdn.openai.com/pdf/67869394-cb91-4c12-888c-5cbd85c7814c/OpenAI-Hugging-Face%20Incident-Technical-Report.pdf

148 **InfoQ.** "AGNTCon + MCPCon 2026 (Agentic AI Foundation adds four projects; MCP security focus)." 26 Oct 2026.
https://www.infoq.com/news/2026/10/agntcon-mcpcon-2026/

149 **Cloudflare.** "Introducing cf: a CLI built for agents." 28 Sep 2026.
https://blog.cloudflare.com/cloudflare-cf-cli-launch

150 **InfoQ.** "Uber's MCP Gateway for AI agents." 2 Nov 2026.
https://infoq.com/news/2026/11/uber-mcp-gateway-ai-agents

151 **King Faisal Specialist Hospital & Research Centre.** "KFSHRC reduces emergency department inpatient bed waiting time by 75%." 15 Sep 2025.
https://www.kfshrc.edu.sa/en/news/2025/09/kfshrc-reduces-emergency-department-inpatient-bed-waiting-time-by

152 **Janke et al., JAMA Network Open.** 2022.
https://pmc.ncbi.nlm.nih.gov/articles/PMC9526134/

153 **Strömblad et al., JAMA Surgery.** 2021.
https://jamanetwork.com/journals/jamasurgery/fullarticle/2775620

154 **Zhang et al., Journal of the American Medical Informatics Association.** 2021.
https://pmc.ncbi.nlm.nih.gov/articles/PMC8633668/

155 **Levin et al., BMJ Innovations.** 2021.
https://doi.org/10.1136/bmjinnov-2020-000420

156 **Kantheti et al., JAMA Network Open.** 3 Sep 2026.
https://pmc.ncbi.nlm.nih.gov/articles/PMC13543018/

157 **Novara Media.** "Hospitals see no notable improvement after using Palantir tool." 28 Jul 2026.
https://novaramedia.com/2026/07/28/hospitals-see-no-notable-improvement-after-using-palantir-tool/

158 **Wei et al., Communications Medicine.** 2024.
https://pmc.ncbi.nlm.nih.gov/articles/PMC11574281/

159 **BNN Bloomberg (CTV News).** "Could AI help reduce surgical backlog? Manitoba researchers are 'very optimistic'." 31 Oct 2026.
https://www.bnnbloomberg.ca/business/artificial-intelligence/2026/10/31/could-ai-help-reduce-surgical-backlog-manitoba-researchers-are-very-optimistic/

160 **Becker's Hospital Review.** "Larry Ellison's AI pitch to Epic hospitals: maybe ours is better." 28 Oct 2026.
https://beckershospitalreview.com/healthcare-information-technology/ehrs/larry-ellisons-ai-pitch-to-epic-hospitals-maybe-ours-is-better

161 **InterSystems.** "Medcare - First Healthcare Provider in EMEA to Adopt InterSystems IntelliCare, a Next-Generation native AI-First Electronic Health Record Platform." 25 Jun 2026.
https://www.intersystems.com/CareBetter/news/medcare-first-healthcare-provider-in-emea-to-adopt-intersystems-intellicare-a-next-generation-native-ai-first-electronic-health-record-platform/

162 **HIT Consultant.** "HIMSS26: Greenway Health Launches Novare, an Agentic AI-Powered EHR Platform for Ambulatory Practices." 10 Mar 2026.
https://hitconsultant.net/2026/03/10/greenway-health-launches-novare-agentic-ai-ehr-ambulatory-care

163 **Fierce Healthcare.** "Enzo Health launches agentic EHR for home health agencies." Jun 2026.
https://www.fiercehealthcare.com/ai-and-machine-learning/enzo-health-launches-ai-driven-enzo-ehr-first-agentic-system-home-health

164 **UrgentIQ press release (PR Newswire, via Webull).** "UrgentIQ Announces Strategic Growth Investment From Five Elms Capital." 11 Sep 2026.
https://www.webull.com/news/15560039573611520

165 **National Informatics Centre (NIC), MeitY.** "NextGen eHospital." 27 May 2026.
https://xn--m1bet4hqd2b.xn--11b7cb3a6a.xn--h2brj9c/project/nextgen-ehospital/

166 **Digital Health News.** "Tamil Nadu's Digital Health Push: Strengthening the State's Public Health System Through Connected Care." Oct 2026.
https://www.digitalhealthnews.com/tamil-nadu-s-digital-health-push-strengthening-the-state-s-public-health-system-through-connected-care

167 **PR Newswire.** "Healthcare Triangle Subsidiary QuantumNexis Inc. Launches its Dubai, UAE Operations to accelerate Digital Health Transformation across the GCC Countries." 27 Feb 2026.
https://stocktitan.net/news/HCTI/healthcare-triangle-subsidiary-quantum-nexis-inc-launches-its-dubai-k45em9s46l8v.html

168 **ThePrint (ANI press release).** "Surgy Health launches SurgyOne, an agentic AI-first HIS/EHR for mid-market hospitals in India and the GCC." Oct 2026.
https://theprint.in/ani-press-releases/surgy-health-launches-surgyone-an-agentic-ai-first-his-ehr-for-mid-market-hospitals-in-india-and-the-gcc/3079612/

169 **BioSpectrum India.** "Prodoc AI secures seed funding to develop Intelligence Layer for Indian hospitals." 30 Sep 2026.
https://www.biospectrumindia.com/news/102/28575/prodoc-ai-secures-seed-funding-to-develop-intelligence-layer-for-indian-hospitals.html

170 **Press Information Bureau / MeitY.** "Digital Personal Data Protection Rules, 2025 (factsheet)." Nov 2025.
https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc20251117695301.pdf

171 **KUOW.** "Records show confusion, finger-pointing during AI Medicare rollout." 27 Oct 2026.
https://www.kuow.org/health-care/2026-10-27/records-show-confusion-finger-pointing-during-ai-medicare-rollout

172 **CSO Online.** "Security spending is growing, except for the typical CISO." 17 Sep 2026.
https://www.csoonline.com/article/4223455/security-spending-is-growing-except-for-the-typical-ciso.html

173 **TechTarget.** "Cyber budgets flatline as CISOs shift new spending to AI." 16 Sep 2026.
https://www.techtarget.com/cybersecurity/news/366650512/Cyber-budgets-flatline-as-CISOs-shift-new-spending-to-AI

174 **Google Cloud.** "Strong vital signs: Healthcare and life sciences ready for AI innovation." 16 Oct 2025.
https://cloud.google.com/transform/healthcare-and-life-sciences-ai-innovation-gen-ai-agents

175 **The Next Web.** "AI scribes are getting drug names and diagnoses wrong in NHS records." 31 Aug 2026.
https://thenextweb.com/news/nhs-ai-scribes-errors-healthwatch-mhra

176 **TechCrunch.** "Medical records giant Epic pauses product development to fix security bugs that risk patients' data." 2 Oct 2026.
https://techcrunch.com/?p=3173140

177 **Becker's Hospital Review.** "Epic's cybersecurity push: 6 things to know." 5 Oct 2026.
https://www.beckershospitalreview.com/healthcare-information-technology/ehrs/epics-cybersecurity-push-6-things-to-know/

178 **The HIPAA Journal.** "Democratic Senators Reintroduce the Health Infrastructure Security and Accountability Act." 18 Sep 2026.
https://hipaajournal.com/health-infrastructure-security-and-accountability-act-2026

APPENDIX C · REFERENCES (CONTINUED)

179 **MPR News.** "Doctors strike, rally for contract with Allina." 14 Sep 2026. https://www.mprnews.org/story/2026/09/14/alina-doctors-strike-rally-for-contract-in-minnesota

180 **Healthcare IT News.** "California Gov. Newsom signs 3 bills to govern healthcare AI." Undated; accessed Nov 2026. https://www.healthcareitnews.com/news/california-gov-newsom-signs-3-bills-govern-healthcare-ai

181 **Axios.** "States take the lead policing AI in health care." 13 Dec 2025. https://www.axios.com/2025/12/13/states-ai-health-care-policy-laws

182 **McGuireWoods.** "HHS finalizes HTI-5, narrowing health IT certification and deferring information-blocking changes." 30 Oct 2026. https://www.mcguirewoods.com/client-resources/alerts/2026/10/hhs-finalizes-hti-5-narrowing-health-it-certification-and-deferring-information-blocking-changes/

183 **Becker's Hospital Review.** "Federal officials weigh how Medicare could pay for 'AI doctors'." 14 Sep 2026. https://www.beckershospitalreview.com/healthcare-information-technology/ai/federal-officials-weigh-how-medicare-could-pay-for-ai-doctors/

184 **HFMA.** "CMS ACCESS model expansion." 21 Sep 2026. https://www.hfma.org/payment-reimbursement-and-managed-care/cms-access-model-expansion/

185 **Press Information Bureau (PIB).** "Ayushman Bharat Digital Mission: India's Digital Health Backbone." 6 Jul 2026. https://static.pib.gov.in/WriteReadData/specificdocs/documents/2026/jul/doc202676912801.pdf

186 **Department of Health - Abu Dhabi.** "Responsible Artificial Intelligence (AI) Standard (DoH/ST/DDGO/RAI/V1/2025)." Oct 2025. https://www.doh.gov.ae/-/media/B84FF5DD027846699C51F36C439A580F.ashx

187 **Middle East AI News.** "UAE Cabinet approves AI services." 18 May 2026. https://www.middleeastainews.com/p/uae-cabinet-approves-ai-services

188 **Inc42.** "CERT-In Flags Unauthorised AI Agent Access To Indian Organisations' Systems." 16 Oct 2026. https://inc42.com/buzz/cert-in-flags-unauthorised-ai-agent-access-to-indian-organisations-systems/

189 **American Hospital Association.** "AHA details recommendations to FDA on regulatory frameworks for GenAI-enabled medical devices." 19 Oct 2026. https://www.aha.org/news/headline/2026-10-19-aha-details-recommendations-fda-regulatory-frameworks-genai-enabled-medical-devices

190 **American Medical Association.** "AMA advances framework for clinical AI liability." 29 Oct 2026. https://www.ama-assn.org/press-center/ama-press-releases/ama-advances-framework-clinical-ai-liability

191 **The Next Web.** "Parents sue OpenAI over ChatGPT advice to cut daughter's seizure drug." 28 Oct 2026. https://thenextweb.com/news/parents-sue-openai-chatgpt-seizure-medication

192 **Becker's Hospital Review.** "What health system CIOs won't automate with AI." 22 Oct 2026. https://www.beckershospitalreview.com/healthcare-information-technology/digital-health/what-health-system-cios-wont-automate-with-ai/

193 **Becker's Hospital Review.** "AHA, Anthropic help shape state oversight of autonomous clinical AI." 14 Oct 2026. https://www.beckershospitalreview.com/healthcare-information-technology/ai/aha-anthropic-help-shape-state-oversight-of-autonomous-clinical-ai/

194 **Newsweek.** "AI could push high-risk ER patients down the queue, study warns." Oct 2026. https://www.newsweek.com/ai-could-push-high-risk-er-patients-down-the-queue-study-warns-12599191

195 **American Medical Association.** "Patients don't have right to ambient AI recordings, court says." 10 Sep 2026. https://www.ama-assn.org/practice-management/digital-health/patients-don-t-have-right-ambient-ai-recordings-court-says

196 **Becker's Hospital Review.** "OpenAI model breached government site to pull hospital bed data." 16 Oct 2026. https://www.beckershospitalreview.com/healthcare-information-technology/ai/openai-model-breached-government-site-to-pull-hospital-bed-data/

197 **Tech Times.** "Microsoft 2026 security report: autonomous ransomware has hacked real organizations." 2 Oct 2026. https://www.techtimes.com/articles/328467/20261002/microsoft-2026-security-report-autonomous-ransomware-has-hacked-real-organizations.htm

198 **Information Age (ACS).** "OpenAI agent hacks Medicare web portal." 24 Sep 2026. https://ia.acs.org.au/article/2026/openai-agent-hacks-medicare-web-portal.html

199 **TechRadar.** "'I broke something': a Claude Code AI agent deleted 48,000 files in just over 100 seconds." 27 Sep 2026. https://techradar.com/pro/security/i-broke-something-a-claude-code-ai-agent-deleted-48-000-files-in-just-over-100-seconds-then-apologized-for-doing-so

200 **Computing.** "Advanced AI models targeted real people during safety tests, UK institute says." Oct 2026. https://www.computing.co.uk/news/2026/ai/advanced-ai-models-targeted-real-people-during-safety-tests-uk-institute-says

201 **The Hacker News.** "Malicious file in Downloads folder hijacks AI coding agents." 26 Oct 2026. https://thehackernews.com/2026/10/malicious-file-in-downloads-folder.html

202 **Cloudflare.** "KV cache exposure on Workers AI: remediation report." 26 Oct 2026. https://blog.cloudflare.com/kv-cache-exposure-remediation/

203 **Let's Data Science.** "Research benchmark examines coding agents' out-of-scope actions." 19 Oct 2026. https://letsdatascience.com/news/research-benchmark-examines-coding-agents-out-of-scope-actio-7fd5b87f

204 **Infosecurity Magazine.** "Billion-dollar companies twice as likely to leave critical vulnerabilities unpatched." 14 Oct 2026. https://www.infosecurity-magazine.com/news/companies-leave-critical/

205 **Becker's Hospital Review.** "Anthropic opens Mythos-class cyber tools to regional hospitals." 7 Oct 2026. https://www.beckershospitalreview.com/healthcare-information-technology/ai/anthropic-opens-mythos-class-cyber-tools-to-regional-hospitals/

206 **CNCF.** "Filtering and reducing AI agent telemetry with the OpenTelemetry Collector." 22 Oct 2026. https://www.cncf.io/blog/2026/10/22/filtering-and-reducing-ai-agent-telemetry-with-the-opentelemetry-collector/

207 **LangChain.** "The components of building a company harness." 24 Oct 2026. https://www.langchain.com/blog/components-building-company-harness

208 **Forkast.** "Schwab and JPMorgan block unauthorized AI agents from moving money." 17 Oct 2026. https://forkast.news/schwab-and-jpmorgan-block-unauthorized-ai-agents-from-moving-money-as-private-ordering-reaches-financial-core/

209 **Techzine.** "Google Agent Compute Engine provides each AI agent with its own microVM." 23 Oct 2026. https://www.techzine.eu/news/security/145226/google-agent-compute-engine-provides-each-ai-agent-with-its-own-microvm/

210 **VKTR.** "Claude Duet: Anthropic previews an AI agent that messages you first." 22 Oct 2026. https://www.vktr.com/ai-platforms/claude-duet-anthropic-previews-an-ai-agent-that-messages-you-first/

211 **NVIDIA.** "NVIDIA launches Open Agent Safety Platform." 28 Sep 2026. https://nvidianews.nvidia.com/news/open-agent-safety-platform

212 **InfoQ.** "AWS uses AI agents to triage Linux kernel CVEs." 27 Oct 2026. https://www.infoq.com/news/2026/10/aws-ai-agents-linux-kernel-cves/

213 **Product Growth.** "Even Healthcare raises $50 mn to expand Bengaluru hospital network." 20 Oct 2026. https://productgrowth.in/news/2026-10-20/even-healthcare-raises-50-mn-to-expand-bengaluru-hospital-network/

214 **Healthcare Dive.** "Suki raises $90M in Series E round to create AI agents." 13 Oct 2026. https://www.healthcaredive.com/news/suki-raises-90m-in-series-e-round-to-create-ai-agents/832862/

215 **GlobeNewswire (Baba).** "Applied AI Company Baba Raises $40 Million Series A to Improve the Medicare Experience." 20 Oct 2026. https://www.globenewswire.com/news-release/2026/10/20/3384598/0/en/applied-ai-company-baba-raises-40-million-series-a-to-improve-the-medicare-experience.html

APPENDIX C · REFERENCES (CONTINUED)

216 **Inc42.** "Clinics On Cloud Bags ₹120 Cr To Take Its Health ATMs Global." 3 Nov 2026.
https://inc42.com/buzz/clinics-on-cloud-bags-₹120-cr-to-take-its-health-atms-global/

217 **Kyndryl press release (PR Newswire, via Finviz).** "Kyndryl completes acquisition of Healthcare IT Leaders, LLC." 14 Sep 2026.
https://finviz.com/news/391401/kyndryl-completes-acquisition-of-healthcare-it-leaders-llc

218 **Neuberger Berman.** "Neuberger and KKR to Invest in Datavant Alongside New Mountain Capital." 15 Sep 2026.
https://www.nb.com/who-we-are/newsroom/neuberger-and-kkr-to-invest-in-datavant-alongside-new-mountain-capital

219 **FemTech Insider.** "Elation Health Acquires Women's Health EHR Aster to Build Agentic Primary Care Platform." 5 Jun 2026.
https://femtechinsider.com/elation-health-acquires-womens-health-ehr-aster-to-build-agentic-primary-care-platform/

220 **Enlivio Health (formerly EMIS) newsroom.** "Enlivio: A bold new chapter begins." 11 Aug 2026.
https://www.enliviohealth.com/news-insights/enlivio-a-bold-new-chapter-begins

221 **BioSpectrum India.** "Cent raises $15 M in Series A to build full-stack preventive health platform for India." 27 Oct 2026.
https://biospectrumindia.com/news/16/28729/cent-raises-15-m-in-series-a-round-to-build-full-stack-preventive-health-platform-for-india.html

222 **Elets eHealth.** "CureBay raises ₹100 crore in pre-Series C led by IvyCap." 28 Oct 2026.
https://ehealth.eletsonline.com/2026/10/curebay-raises-100-crore-pre-series-c-ivycap-ai-rural-healthcare-andhra-pradesh

223 **Anthropic.** "Claude Frontier Academy." 2 Oct 2026.
https://www.anthropic.com/news/claude-frontier-academy

224 **Bain & Company.** "AI in Indian Healthcare Delivery." Sep 2026.
https://www.bain.com/de/insights/ai-in-indian-healthcare-delivery/

225 **Digital Health News.** "NHA's new guidebook maps out what a model digital health facility looks like, step by step." 6 Oct 2026.
https://www.digitalhealthnews.com/nha-s-new-guidebook-maps-out-what-a-model-digital-health-facility-looks-like-step-by-step

226 **National Accreditation Board for Hospitals & Healthcare Providers (NABH).** "Draft NABH Standards for HIS and EMR Systems - 23rd July 2024." 23 Jul 2024.
https://portal.nabh.co/Announcement/Draft%20NABH%20Standards%20for%20HIS%20and%20EMR%20Systems%20-%2023rd%20July%202024.pdf

227 **Healthcare IT News (Asia).** "India to harness ABDM data to develop public AI." Sep 2024.
https://www.healthcareitnews.com/news/asia/india-harness-abdm-data-develop-public-ai

228 **ThePrint.** "NHA, IIT Kanpur partner to develop open platform for comparing AI models." 11 Sep 2024.
https://theprint.in/india/nha-iit-kanpur-partner-to-develop-open-platform-for-comparing-ai-models/2263388

229 **Digital Health News.** "NHA CEO Meets UAE Health Officials to Explore Digital Health Partnership." 26 Oct 2026.
https://www.digitalhealthnews.com/nha-explores-collaboration-with-uae-on-digital-health-healthcare-ai

230 **Medical Dialogues.** "200 Private Hospitals face action for non-compliance under Ayushman Bharat in Uttar Pradesh." 16 May 2026.
https://medicaldialogues.in/amp/news/health/hospital-diagnostics/200-private-hospitals-face-action-for-non-compliance-under-ayushman-bharatin-uttar-pradesh-170817

231 **Press Information Bureau.** "Press release on ABDM, eSanjeevani and AB-PMJAY progress (PRID 2313393)." 17 Sep 2026.
https://www.pib.gov.in/PressReleasePage.aspx?PRID=2313393&reg=48&lang=2

232 **Press Information Bureau.** "Factsheet on digital health (February 2026)." Feb 2026.
https://static.pib.gov.in/WriteReadData/specificdocs/documents/2026/feb/doc2026213788701.pdf

233 **Press Information Bureau / National Health Authority.** "Press release on PM-JAY claims processing (PRID 2251434)." Apr 2026.
https://www.pib.gov.in/PressReleasePage.aspx?PRID=2251434&reg=3&lang=1

234 **NHSRC.** "National Health Accounts estimates for India, 2022–23." May 2026.
https://nhsrcindia.org/sites/default/files/2026-05/NHA%202022-23%20Report.pdf

235 **Press Information Bureau.** "Press release on doctor and nurse availability (PRID 2225755)." Feb 2026.
https://www.pib.gov.in/PressReleasePage.aspx?PRID=2225755&reg=3&lang=1

236 **Press Information Bureau.** "IndiaAI Mission factsheet." Feb 2026.
https://static.pib.gov.in/WriteReadData/specificdocs/documents/2026/feb/doc2026217795001.pdf

237 **Saudi Ministry of Health.** "News release, 9 August 2026." 9 Aug 2026.
https://www.moh.gov.sa/en/ministry/mediacenter/news/pages/news-2026-08-09-001.aspx

238 **Forbes Middle East.** "Top Healthcare Leaders 2026: Founders & Shareholders." 16 Sep 2026.
https://www.forbesmiddleeast.com/lists/top-healthcare-leaders-2026-founders-shareholders/

239 **Sharikat Mubasher (Argaam).** "Over 2K brands to participate in Global Health Exhibition next October." 12 Aug 2026.
https://sharikatmubasher.com/news/article/21568180/over-2k-brands-to-participate-in-global-health-exhibition-next-october?lang=en

240 **Healthcare IT News (EMEA).** "Understanding NPHIES, KSA's health information exchange." 2023.
https://www.healthcareitnews.com/news/emea/understanding-nphies-ksas-health-information-exchange

241 **Computer Weekly.** "Saudi Arabia launches AI health coach to transform preventive care." Oct 2025.
https://computerweekly.com/news/366633677/Saudi-Arabia-launches-AI-health-coach-to-transform-preventive-care

242 **Middle East AI News.** "Saudi to launch world's largest AI physician clinical trial." 5 Nov 2025.
https://www.middleeastainews.com/p/saudi-to-launch-worlds-largest-ai

243 **Saudi Food and Drug Authority.** "SFDA grants marketing authorisation to first Saudi-developed AI-enabled medical software." 10 Sep 2026.
https://sfda.gov.sa/en/news/19462

244 **HPE.** "Saudi German Health UAE transforms digital healthcare in the UAE with the HPE Self-Driving Network." 29 Sep 2026.
https://www.hpe.com/us/en/newsroom/press-release/2026/09/saudi-german-health-uae-transforms-digital-healthcare-in-the-uae-with-the-hpe-self-driving-network.html

245 **Saudi Ministry of Health.** "News release, 27 January 2026." 27 Jan 2026.
https://www.moh.gov.sa/en/ministry/mediacenter/news/pages/news-2026-01-27-002.aspx

246 **Malaffi.** "Malaffi reaches 3.5 billion clinical records." 2025.
https://www.malaffi.ae/malaffi-reaches-3-5-billion-clinical-records-setting-a-new-standard-for-healthcare-innovation-in-abu-dhabi/

247 **UAE Ministry of Health and Prevention.** "About Riayati." 2026.
https://mohap.gov.ae/en/riayati/about-numr

248 **General Authority for Statistics, Saudi Arabia.** "Healthcare establishments and workforce statistics 2023." 2023.
https://www.stats.gov.sa/documents/d/guest/healthcare-establishments-and-workforce-statistics-2023en-pdf

249 **CERT-In.** "Directions under sub-section (6) of section 70B of the IT Act, 2000." 28 Apr 2022.
https://www.cert-in.org.in/PDF/CERT-In_Directions_70B_28.04.2022.pdf

250 **Ministry of Electronics and Information Technology.** "Digital Personal Data Protection Rules, 2025 (G.S.R. 846(E))." 13 Nov 2025.
https://www.meity.gov.in/static/uploads/2025/11/53450e6e5dc0bfa85ebd78686cadad39.pdf

251 **Ministry of Electronics and Information Technology.** "Digital Personal Data Protection Act, 2023." 11 Aug 2023.
https://www.meity.gov.in/static/uploads/2024/06/2bf1f0e9f04e6fb4f8fef35e82c42aa5.pdf

252 **SCC Online.** "DPDP (Removal of Difficulties) Order, 2026." 15 Oct 2026.
https://www.scconline.com/blog/post/2026/10/15/dpdp-removal-of-difficulties-order-2026/

APPENDIX C · REFERENCES (CONTINUED)

253 **Surgy Health.** "DPDP Act for hospitals." Undated; accessed Nov 2026. https://surgy.health/blog/dpdp-act-for-hospitals

254 **NABH.** "Hospital Accreditation Standards, 6th Edition." Jan 2025. https://portal.nabh.co/images/Standards/NABH%20Hospital%20Accreditation%20Standard%206th%20Edition%20January%202025.pdf

255 **Digital Health News.** "NABH introduces 6th edition of accreditation standards for hospitals with emphasis on digital health technology." 29 Nov 2024. https://www.digitalhealthnews.com/nabh-introduces-6th-edition-of-accreditation-standards-for-hospitals-with-emphasis-on-digital-health-technology

256 **Surgy Health.** "NABH 6th edition roadmap 2026." Undated; accessed Nov 2026. https://surgy.health/blog/nabh-6th-edition-roadmap-2026

257 **NABH.** "Digital Health Standards for Hospitals, 2nd Edition." Sep 2025. https://nabh-portal-live.s3.ap-south-1.amazonaws.com/wp-content/uploads/2025/07/02133924/NABH-Digital-Health-Standards-For-Hospitals-2nd-Edition_September-2025.pdf

258 **Surgy Health.** "Hospital KPI dashboard for NABH and JCI." 8 Nov 2026. https://surgy.health/blog/hospital-kpi-dashboard-nabh-jci

259 **Surgy Health.** "CDSCO software as a medical device in India." 8 Nov 2026. https://surgy.health/blog/cdsco-software-as-medical-device-india

260 **Emergo by UL.** "India CDSCO finalizes guidance on medical device software." 30 Jul 2026. https://www.emergobyul.com/news/india-cdsco-finalizes-guidance-medical-device-software

261 **Surgy Health.** "AI in healthcare in the UAE and Saudi Arabia." 8 Nov 2026. https://surgy.health/blog/ai-in-healthcare-uae-saudi-arabia

262 **Saudi Food and Drug Authority.** "Guidance for AI/ML-enabled medical devices (MDS-G010)." 3 Jan 2023. https://sfda.gov.sa/en/guide/guidance-artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices-mds-%E2%80%93-g010

263 **Polsinelli.** "TJC and CHAI release guidance on responsible AI use in healthcare." Sep 2025. https://www.polsinelli.com/publications/tjc-chai-release-guidance-responsible-ai-use-healthcare

264 **Surgy Health.** "Hospital AI governance policy template." 8 Nov 2026. https://surgy.health/blog/hospital-ai-governance-policy-template

265 **Surgy Health.** "Healthcare AI RFP template and scorecard." 8 Nov 2026. https://surgy.health/blog/healthcare-ai-rfp-template-scorecard

266 **Surgy Health.** "Hospital AI buyer's guide." Undated; accessed Nov 2026. https://surgy.health/guides/hospital-ai-buyers-guide

267 **arXiv.** "Clinician use of language models diverges from how the models are evaluated." 8 Oct 2026. https://arxiv.org/abs/2610.11069

268 **TechTarget.** "AI-based clinical intelligence is all the rage. But who is using it?." 19 Oct 2026. https://www.techtarget.com/healthtechanalytics/news/366652436/AI-based-clinical-intelligence-is-all-the-rage-But-who-is-using-it

269 **Corti.** "Symphony for Text Generation: benchmarking clinical note quality." 14 Oct 2026. https://www.corti.ai/stories/symphony-for-text-generation-benchmark-clinical-note-quality

270 **LangChain.** "Common questions about agent evals." 15 Oct 2026. https://www.langchain.com/blog/common-questions-about-agent-evals

271 **arXiv.** "Reliable AI-assisted decisions via selective auditing." 22 Oct 2026. https://arxiv.org/abs/2610.31662

272 **Anthropic.** "An update on independent evaluators." 15 Oct 2026. https://www.anthropic.com/news/an-update-on-independent-evaluators

273 **Anthropic.** "Early progress in drug development." 27 Oct 2026. https://www.anthropic.com/research/early-progress-drug-development

274 **LlamaIndex.** "Introducing the Document Parsing Protocol." 28 Oct 2026. https://www.llamaindex.ai/blog/introducing-the-document-parsing-protocol

275 **TechCrunch.** "AI delivers zero financial payoff to most companies, according to Apollo economist." 25 Oct 2026. https://techcrunch.com/2026/10/25/ai-delivers-zero-financial-payoff-to-most-companies-according-to-apollo-economist/

276 **Surgy Health.** "SurgyOne: agentic AI hospital information system." Undated; accessed Nov 2026. https://surgy.health/products/surgy-one

277 **Surgy Health.** "SurgyScribe: AI medical scribe." Undated; accessed Nov 2026. https://surgy.health/products/surgy-scribe

278 **Surgy Health.** "SurgyFrontdesk: AI voice agent for hospitals." Undated; accessed Nov 2026. https://surgy.health/products/surgy-frontdesk

279 **Surgy Health.** "SurgyCRM: hospital CRM." Undated; accessed Nov 2026. https://surgy.health/products/surgy-crm

280 **Surgy Health.** "Surgy Health: AI-first products for hospitals." Undated; accessed Nov 2026. https://surgy.health/

281 **Surgy Health.** "SurgyInsights: conversational hospital analytics." Undated; accessed Nov 2026. https://surgy.health/products/surgy-insight

282 **Surgy Health.** "SurgyLearn: hospital staff training platform." Undated; accessed Nov 2026. https://surgy.health/products/surgy-lms

283 **Surgy Health.** "AI medical scribe software for hospitals." Undated; accessed Nov 2026. https://surgy.health/blog/ai-medical-scribe-software-for-hospitals

ABOUT THE AUTHORS

The people behind *the book*



Mohammed Jamil Nasir

AUTHOR · FOUNDER – PRODUCT & TECHNOLOGY, SURGY HEALTH

Nasir has more than 14 years of experience across software engineering, product management and UX, including more than five years building AI and generative-AI products for healthcare. He holds a Global MBA from SP Jain, a PG Certificate in AI/ML from IIT Guwahati and a bachelor's degree in computer science engineering, and has worked in Singapore. He leads product and technology at Surgy Health.



Zuhaib Ahmed Khan

CO-AUTHOR · SURGY HEALTH

Zuhaib works on platform engineering, deep tech and agentic AI software development at Surgy Health, building the systems that let AI agents run safely inside hospital software. His daily Platform & Agentic AI Insights kept this book current through early November 2026.



Simran Nagekar

CO-AUTHOR · PRODUCT ENGINEER, SURGY HEALTH

Simran is a product engineer at Surgy Health. Her research notes on EMR/EHR, AI and agentic AI started the research log behind this book, and she wrote Surgy Health's guide to hospital AI budgets for 2027, on which Chapter 6 draws.

SURGY. health

ABOUT THE PUBLISHER

Surgy Health is an AI-first healthcare technology company based in Bengaluru, India, building seven products for hospitals in India and the GCC: **SurgyOne**, an agentic-AI hospital information system, with **SurgyScribe**, **SurgyFrontdesk**, **SurgyCRM**, **SurgySettle**, **SurgyInsights** and **SurgyLearn**. The next two pages describe them. Surgy Health is a brand of Surgy Innovation Labs Private Limited.

WEB surgy.health EMAIL info@surgyy.com OFFICE J. P. Nagar, Bengaluru 560078

ABOUT THE PUBLISHER

What Surgy Health *builds*

The chapters of this book describe categories, not our products. These two pages do the opposite: they describe what Surgy Health builds for hospitals in India and the GCC, in our own words. Read them as a vendor's description, and test it as Chapters 7, 13 and 14 suggest.

SurgyOne HOSPITAL INFORMATION SYSTEM

An agentic-AI hospital information system with one patient record across OPD, IPD, emergency, labs, pharmacy and billing. AI agents inside each module prioritise ward rounds, flag claim risks and prepare discharge summaries for staff to approve; every AI action is logged and approved by the hospital's team. Built on a FHIR R4 data model, designed for ABDM milestones M1–M4 and NHCX, with NPHIES among its listed integrations.[276]

Phased roll-out, starting with OPD, labs and billing · tenant-isolated data, role- and branch-scoped access, full audit trail[276]

SurgyScribe CLINICAL DOCUMENTATION

Turns consultations into structured notes from department templates, suggests ICD-10 and ICD-11 codes and writes finalised notes back to the EHR. Handles mixed-language consultations; the doctor remains accountable for every note.[277]

Patient identifiers are not sent to third-party AI services; AI provider accounts configured for zero data retention[277]

SurgyFrontdesk VOICE AGENT

Answers and makes patient calls on the hospital's existing phone line: booking, rescheduling, registration, common questions and follow-up calls, with bills and reports sent by WhatsApp or SMS.[278]

Out-of-scope and complex calls go to staff; calls with negative sentiment are flagged[278]

SurgyCRM PATIENT JOURNEYS

One inbox for enquiries from web, WhatsApp, call centre, walk-ins, campaigns and referring doctors, with milestone-based patient journeys and bedside feedback. A lead is marked verified only when it matches a real billing event in the HIS.[279]

WhatsApp-first, through the hospital's own verified business account[279]

SurgySettle CLAIMS AND RECEIVABLES

Automates discharge summaries, claims, TPA work and receivables follow-up for the revenue cycle.[280]

Connects to SurgyOne's billing and claims workflow[276,280]

SurgyInsights ANALYTICS

Plain-language questions over hospital data in English, Hindi or Arabic, with 84 KPIs across seven departments tagged to NABH and JCI, alerts and root-cause analysis. A language model interprets the question; a deterministic formula engine computes every number.[281]

Starts from CSV or Excel exports; personal identifiers blocked at upload[281]

SurgyLearn STAFF TRAINING

Turns a hospital's own SOPs into narrated, quizzed courses, assigns them by role through HRMS sync, proctors assessments and exports training records for NABH and JCI audits.[282]

24 ready-made hospital courses, from BLS to hand hygiene[282]

ABOUT THE PUBLISHER (CONTINUED)

In practice: results our customers report

More than 4,000 healthcare professionals across more than 500 centres use Surgy Health products.[266,280] The figures below are reported by Surgy Health and its customers. They have not been independently audited, and we would hold any other vendor's numbers to the same caveat.

SURGYSCRIBE · ONE HOSPITAL, ONE MONTH

60%

of 291 AI-drafted note sections accepted without edits; 97% average section accuracy; 13 seconds average processing per note.[266,283]

SURGYSCRIBE · APOLLO CLINIC, KUWAIT

10–15 → 1–1.5 min

documentation time per note, in English and Arabic, across a multispecialty clinic of six to eight doctors.[266,283]

SURGYCRM · MEDIVERSAL HOSPITALS

600+

calls handled a day across three hospitals of 100–150 beds each, with more than 200 patient-journey milestones a week.[279]

SURGYLEARN · DCDC KIDNEY CARE

2,687

staff onboarded across 30 departments since July 2025; 51,093 training assignments delivered, 99% of them automatically; 819 course completions.[282]

SURGYLEARN · AL-SALAM HOSPITALS

A phased roll-out from August 2026: 98 users in phase one, scaling to more than 1,000, with 19 courses live.[282]

How we build

AI drafts, people decide
Every AI action that touches a record, a claim or a patient waits for staff approval and is logged.[276]

One record
SurgyOne keeps OPD, IPD, emergency, labs, pharmacy and billing on one patient record, on FHIR R4.[276]

Start from exports
Analytics and CRM can start from CSV or Excel exports while HIS integration is built.[279,281]

Each hospital its own tenant
Data isolated per hospital, encrypted in transit and at rest, with role- and branch-scoped access.[280]

National rails
Designed for ABDM, NHCX and NPHIES, and for India's DPDP Act.[276,280]

Three languages
English, Hindi and Arabic for staff and patients; the scribe handles mixed-language consultations.[277,280]

CONTACT

surgy.health · info@surgyy.com · Surgy Innovation Labs Private Limited, J. P. Nagar, Bengaluru 560078, India. Security and privacy questions: nasir@surgyy.com.

Your next hospital system won't just record care. *It will help run it.*

Hospital AI is moving through three stages: listening, acting and orchestrating. Drawing on 56 research briefs and daily digests, the latest adoption and outcome data, and fact-checked news from the US, India and the GCC through November 2026, this field guide shows hospital CEOs, CIOs, medical directors and finance heads what changed in 2026, what is working, what to avoid, and how to choose and govern an AI-first hospital information system in 2027.

- Why the EHR is becoming a platform that AI agents plug into
- Where agentic AI already pays back, and how to prove it on your own data
- How to keep AI agents inside their limits, from identity to audit
- ABDM, DPDP, CERT-In, NPHIES and the Gulf's new AI rules
- A buyer's scorecard, a 90-day plan and a 24-point readiness check



Mohammed Jamil Nasir
FOUNDER – PRODUCT & TECHNOLOGY



Zuhaib Ahmed Khan
PLATFORM ENGINEERING & AGENTIC AI



Simran Nagekar
PRODUCT ENGINEERING & RESEARCH

SURGY. health

HEALTHCARE MANAGEMENT · HEALTH INFORMATICS
ARTIFICIAL INTELLIGENCE · SURGY.HEALTH

SCAN · SURGY.HEALTH

ISBN 978-93-346-0290-6
9 789334 602906