

PERSPECTIVES ON CONNECTED CARE / 2026

The Cost of Standing Still

When does adding intelligence to yesterday's healthcare stack cost more than changing it?

An economic and architectural argument for spending intelligence where it changes the outcome, and for knowing, rather than assuming, what the current stack costs.

Echo Health AI

Visualisation Hub Pty Ltd trading as Echo Health AI

Public edition | 29 September 2026

Australian focus. An economic and design argument, not an evaluation of any product. Vendor descriptions are dated observations of public material and are labelled as vendor claims.

In this paper

THE QUESTION

Executive summary	03
1. From scribe to agent	04
2. Agency has an observable price	05

THE CASE FOR THE STATUS QUO

3. Why healthcare stands still	06
4. The accumulation problem	07

THE AGENTIC RESPONSE

5. Agents are a rational response	08
6. The market is splitting	09

THE ECONOMICS OF INTELLIGENCE

7. Agency is not computationally free	10
8. Full cost per correct, verified outcome	11
9. Does this need an agent?	12
10. Three taxes: reasoning, fragmentation, integration	13

WHERE INTELLIGENCE BELONGS

11. Where should healthcare spend intelligence?	14
12. Expanding the clinical frontier	15
13. Intelligence by design	16

THE DECISION

14. The economics of familiarity	17
15. The crossover question	18
16. The burden of proof, and what we do not yet know	19
17. A decision framework, and conclusion	20

SOURCES

References	21
------------	----

About this edition

This paper joins Echo Health AI's 2026 series with *After the Scribe*, *Informed Consent and Shared Benefit* and *The Digital Front Door—and the Forgotten Back Door*. It asks an economic question that the move from generative to agentic AI makes urgent. Products are described from their own public material, dated and labelled as vendor claims; they illustrate architectural strategies and are not ranked. Figures marked conceptual or illustrative contain no data. Sources are numbered in the text and linked in the reference pages.

03 / THE ECONOMICS OF CONNECTED CARE

Executive summary

Healthcare AI is moving from generating content to performing work. On 28 September 2026, Heidi Health announced Heidi II, which plans and carries out administrative work across the record, calendar, inbox and portal and, by default, returns it for review. Where clinicians work across disconnected systems, agents like this can create real value.

This paper does not ask whether agents are good. It asks when adding an intelligent layer to a fragmented stack becomes more expensive than changing the architecture beneath it, and, more usefully, where agentic intelligence is valuable enough to justify its cost.

Keeping an incumbent system is often rational. Replacing a core platform means migration, training, disruption and risk, and the published evidence on record-system transitions shows those costs are real. Adding an agent is incremental and can pay back quickly. But rational additions accumulate: point solutions, integrations and now agent layers, each with a recurring price.

Agency has an observable market price, and a computational one. Neither makes agents poor value. Both mean the right measure is the **full cost per correct, verified outcome**, compared with the real operational baseline rather than with zero.

Many steps in care are not reasoning problems. Whether a result has returned, been reviewed or become overdue is state, time and workflow. Using probabilistic reasoning for a step that could be deterministic creates a **reasoning tax**; distributing work and state across disconnected systems creates **fragmentation** and **integration taxes**. Agents can reduce the last two while adding the first.

Echo Health AI's position is **intelligence by design**: build the patient journey and the operating environment coherently, then use the appropriate mechanism at each step, whether that is an agent, a reasoning model, specialised clinical AI, deterministic software, workflow or human clinical judgement. Some of the most valuable healthcare AI does not compensate for complexity at all. It extends what clinicians can see.

The point at which changing becomes cheaper than staying cannot yet be specified from independent evidence, and a replacement platform carries the higher burden of proof. This paper offers a way to calculate it locally.

The question is not agents or no agents. It is where intelligence changes the outcome, and what the whole system costs to deliver it.

04 / THE ECONOMICS OF CONNECTED CARE

1. From scribe to agent

For three years the most visible AI in clinical practice has been the ambient scribe: software that listens to a consultation and drafts the note. The next phase is different in kind. AI no longer only documents work. Increasingly, it performs it.

Heidi II is a clear contemporary example. Heidi Health, founded in Melbourne, announced it on 28 September 2026 as “a new version of its platform that acts on the administrative work surrounding every patient visit rather than documenting the visit alone”, with rollout from 29 September. [1] Its public material describes:

- **Agents** that plan and execute a task, then return the result for review. [1]
- **Memory** that learns a clinician’s preferences and standards from instructions and corrections. [2]
- **Pre-charting**, completing referral forms, chasing pathology, following up results and drafting recalls. [1, 2]
- **Routines** that run across a patient list, including panel sweeps for patients due for immunisations, blood tests, screening or recalls. [1]
- Work “across your calendar, inbox, and portal”, through new **connectors** to practice systems and through **computer use**, operating software “like a member of the care team”. [2, 1]

Heidi describes graduated control. It “starts in Manual mode”; by default, anything that touches a chart, reaches a patient or leaves the system waits for review, and a clinician can switch a trusted task to Auto mode. Some actions are excluded in every mode: Heidi “will not prescribe, delete a record, or submit a claim”. [2] The launch material describes clinical-safety review before and after release and pre-launch testing with clinicians. [1, 3] These are the vendor’s own statements on launch day. No independent evaluation of the agents had been published when this paper went to press.

Heidi’s press release describes the product as “the connective tissue across clinical tools”. [1] The phrase captures both the promise and the question this paper asks. Connective tissue is valuable because what it connects is in pieces.

When AI can operate the stack, the question shifts from what a model can write to what the whole system should cost.

Evidence boundary. Heidi II is used as a dated example of a market shift, not assessed as a product. Descriptions come from Heidi’s public pages as read on 29 September 2026 and may change.

05 / THE ECONOMICS OF CONNECTED CARE

2. Agency has an observable price

At launch, Heidi's Australian pricing interface listed per-user plans that include agent usage. The figures below were captured by Echo Health AI from the live interface and checked against the published page on 29 September 2026. [4] They are Australian dollars per user, and the page states that they exclude applicable sales taxes.

PLAN AND SEAT	BILLING	PER USER / MONTH	PER USER / YEAR
Clinician, Standard seat	Annual	A\$85	A\$1,020
Clinician, Max seat	Annual	A\$249	A\$2,988
Clinician, Max seat	Monthly	A\$349	A\$4,188 over 12 months
Practice, Standard	Annual	A\$130	A\$1,560
Practice, Max	Annual	A\$399	A\$4,788
Enterprise	Not published	Custom	n/a

The Clinician plan includes "Usage for Agents", a monthly allowance consumed by agents, routines, computer use, memory and connectors; the Max seat includes five times the Standard allowance. [4] Annualised, a Practice Max commitment scales linearly with team size:

PRACTICE MAX USERS	10	20	50	100
Annual commitment, before tax	A\$47,880	A\$95,760	A\$239,400	A\$478,800

Excludes applicable taxes and the practice's other software, integration and staff costs. Arithmetic: A\$399 × 12 months × users.

These numbers do not show that agents are poor value. **A\$399 per user per month is excellent value if the agent creates more than A\$399 a month of incremental value for that user**, net of the review time it creates. At an illustrative loaded cost of A\$150 per clinician hour (an assumption; substitute your own), that is about 2.7 hours a month, or roughly eight minutes per working day.

What the figures do show is that agentic orchestration now has an observable market price. It is a line item that can be compared: with the value it creates, with other ways of producing the same outcome, and with the cost of the architecture it operates.

Price alone is not the question. Price per verified outcome, against the real alternative, is.

Evidence boundary. Pricing is a dated observation of a public page and may change. Nothing here states what any customer pays.

3. Why healthcare stands still

Healthcare organisations are often criticised for keeping old software. The criticism usually misreads the decision. **The incumbent does not need to prove it deserves to stay every morning. The challenger has to prove it is worth changing to.**

Evidence on replacing a record system is thin but consistent. A review by an AMIA working group found only five papers that met its criteria and concluded that transitions are “remarkably expensive, laborious, personnel devouring, and time consuming”, citing financial burden, staffing, data integrity in migration, access to legacy records and patient-safety threats. [5] In four US health systems, every site reduced clinical capacity at go-live; one clinic took a dozen or more superusers off patient care in the first week. [6] After Mayo Clinic changed record systems, patient satisfaction with access fell at all six sites studied and took 9 to 15 months to recover. [7] When one academic system moved to a commercial record, interruptive drug-interaction alerts rose six-fold and acceptance of the most severe alerts fell from 100% to 8.4%: years of local tuning did not survive the move. [8] A year after go-live, US Veterans Affairs clinicians were still reporting poor fit, [9] and Denmark’s national auditor found that Capital Region hospitals struggled to treat as many patients after their new platform was rolled out. [10]

Balance matters. A study of 17 US hospital go-lives found no overall increase in short-term mortality, adverse safety events or readmissions. [11] The case for caution rests on cost, disruption and lost capacity, not on a claim that transitions routinely harm patients.

Markets reflect these costs. US hospital records have consolidated onto two vendors, and researchers judge that reversing consolidation is unlikely “due to high EHR switching costs”. [12, 13] In Australia, two clinical information systems were reported as used by 81% of general practices (a figure based on 2016 data). [14] And people disproportionately keep the status quo, even in consequential decisions. [15]

ADD AN AGENT

- Comparatively incremental
- Familiar workflows retained
- Lower immediate disruption
- Potentially fast return
- Can be trialled one task at a time

REPLACE THE CORE PLATFORM

- Data migration and validation
- Implementation, staffing and training
- Workflow redesign and integration change
- Temporary loss of capacity
- Organisational and clinical risk

Figure 1. The asymmetric decision. The costs of adding are small and near; the costs of replacing are large, near and certain, while its benefits are later and uncertain.

Incremental growth of the stack is rational. That is exactly why its cumulative cost goes unexamined.

07 / THE ECONOMICS OF CONNECTED CARE

4. The accumulation problem

Healthcare organisations rarely decide to build a fragmented architecture. They decide, one reasonable purchase at a time, to solve the problem in front of them: online booking to reduce phone load, forms to stop re-keying, payments, telehealth, a scribe, analytics, a portal and now an agent. Each decision can be rational on its own. The aggregate can still be costly and fragmented.




Figure 2. The accumulation problem (illustrative). Not every practice has every layer, and the order varies. “Possible pairings” is the combinatorial maximum, $n(n-1)/2$, not a count of real interfaces: many systems connect only through the core or a hub. It shows why each addition makes the whole harder to hold together.

Each system brings its own subscription, login, data model, support relationship and integration. The number of possible pairings grows faster than the number of systems. No organisation connects every pair, but every connection it makes must be mapped, tested, monitored and maintained, and every gap it leaves becomes manual work.

The time cost of today’s tools is already large. A US time-and-motion study found nearly two hours of record and desk work for every hour of direct patient time, [16] and event logs showed family physicians spending 5.9 hours of an 11.4-hour day in the record, 85 minutes of it on the inbox. [17] These studies do not isolate fragmentation, but they show where time goes when work and information are spread across screens.

Nobody chose the stack. It accumulated.

5. Agents are a rational response

If clinicians genuinely work across disconnected systems, an agent that understands intent and performs work across those systems can create large value. The agentic-overlay case deserves full credit.

LESS NAVIGATION

Fewer application switches, fewer logins, less searching for the same fact in three places.

LESS DUPLICATION

Information entered once and carried to the form, letter or referral that needs it.

FASTER FOLLOW-UP

Outstanding results chased, recalls drafted and handovers prepared without waiting for a quiet hour.

BETTER RETRIEVAL

Relevant history gathered across sources before the clinician asks for it.

ORCHESTRATION

Multi-step work carried across systems that were never designed to talk to each other.

FEWER DROPPED TASKS

Exceptions surfaced by name, in natural language, instead of lost in a queue.

Heidi's own description of that last point is a genuine design strength: "Nothing is skipped silently, and anything Heidi couldn't finish comes back named with a reason." [2] An overlay is also the quickest route to value that replaces nothing. It meets organisations where they are, uses systems staff already know, and can be adopted and evaluated one task at a time. For a practice whose core system is stable and whose pain is administrative friction, it may be the right choice.

Other companies are taking different routes towards the same work. They are examples, not a ranking, and each description is the vendor's own:

- **Tandem Health** (Sweden) provides documentation, coding and decision support that write into more than 100 existing record systems, and describes a phased 2026 move towards "an AI native operating system for clinics", built "on top of the infrastructure clinics already rely on". Its agent features are marked as coming soon. [18]
- **Abridge** (US) describes "one intelligence layer" across health systems, extending ambient documentation inside the record towards context-aware decision support and work after the visit. [19]
- **Ambience Healthcare** (US) describes "one system of context and action" that reconciles information and "writes the decision back into systems of record". [20]
- **Innovaccer** (US) unifies data from records, claims and other sources so that its agents "share context, hand off work, and learn from one source of truth". [21]

Where the systems cannot be changed, an agent that works across them may be the best available mechanism.

Evidence boundary. Quotations are from each vendor's public pages, read on 29 September 2026. They are not independent evaluations, and inclusion implies no judgement of quality.

09 / THE ECONOMICS OF CONNECTED CARE

6. The market is splitting

Beneath the product names, the market is making three broad architectural bets. They differ in where state lives, where reasoning sits and what the intelligence is asked to do.




Figure 3. Three architectural bets. The categories describe emphasis, not fixed boundaries: overlays are adding connectors and memory, and platforms are adding agents. Echo One is being built as an Architecture C platform.

ARCHITECTURE	EXAMPLES, AS EACH DESCRIBES ITSELF	MATURITY
A. Overlay	Heidi II: agents across record, calendar, inbox and portal via connectors and computer use [1] · Tandem Health: agents announced for phased 2026 rollout [18]	Established vendors; agent features new or in rollout
B. Context + action	Ambience: “one system of context and action” [20] · Innovaccer: data platform with agents on shared context [21] · Abridge: “one intelligence layer” [19]	Established, heavily funded vendors
C. AI-native OS	Greenway Health Novare: “designed with agentic AI at its core” [22] · Canvas Medical: record with SDK and command structure for agents [23] · Elation Health: acquired AI-native Aster, June 2026 [24] · mdhub: “AI workers” in a clinic OS [25] · Ocean Health: “AI-native system of record” [26]	Mixed: established EHR vendors launching new platforms; early-stage companies

Evidence boundary. Classification is Echo Health AI’s reading of each vendor’s public description. Greenway’s published efficiency figures come from a single vendor pilot; other scale and funding statements are vendor or press reports. None of these companies, Echo Health AI included, has published independent comparative evidence of total cost of ownership.

7. Agency is not computationally free

The cost of an agent is more than its tokens. A system that plans, retrieves, calls tools and checks its own work uses more computation and time than a single model call, and it needs more operational scaffolding around it.

The most direct recent evidence is a 2026 *npj Digital Medicine* benchmark of agent systems on clinical decision tasks, including simulated doctor–patient dialogues and difficult medical question sets. [27] The best agent outperformed the best baseline model by 0.5 to 8.9 percentage points, differences the authors report did not reach statistical significance. The agents used more than ten times the tokens of their base model (10 to 100 times, in the authors’ discussion) and frequently doubled or tripled response time. Filtering and prompting reduced hallucinations, but they still affected about 30% of clinical scenarios. The authors caution that idealised benchmarks may overestimate reliability.

This is one benchmark of general-purpose agent frameworks on one open model. It does not show that every agent is inefficient, and newer designs may do better. Its value is narrower: **agency is not computationally free, and its incremental benefit should be measured rather than assumed.**

Other work points the same way. On a coding benchmark, one complex agent architecture cost more than fifty times as much as a simple retry strategy without a significant accuracy advantage, which led its authors to call for cost-controlled evaluation. [28] In MedAgentBench, a simulated FHIR record environment, the best model completed 69.7% of physician-written tasks overall but only 54% of tasks that write to the record. [29] Human review is not free insurance either: when a purported AI suggested a wrong mammography category, radiologists’ accuracy fell sharply, including among very experienced readers. [30]

Costs are also falling. Epoch AI estimates that the price of reaching a fixed benchmark level has fallen between 9 and 900 times a year, depending on the task. [31] Cheaper tokens shrink the inference share of an agent’s cost. They do not remove the cost of evaluating, monitoring, reviewing and correcting its work.




Figure 4. What an agent costs. Token prices affect only the first column. The other two scale with the number and risk of the actions an agent takes.

11 / THE ECONOMICS OF CONNECTED CARE

8. Full cost per correct, verified outcome

After the Scribe argued that buyers should measure the complete cost of a verified outcome. [32] Stated as a ratio, that argument becomes the economic core of this paper. It is architecture-neutral: it applies equally to an agent, a deterministic workflow, a specialised model, a human-only process and a replacement platform.




Figure 5. Full cost per correct, verified outcome. The numerator is every cost the workflow incurs; the denominator counts only outcomes that are complete and verified.

The denominator is the discipline. An agent that attempts 2,000 tasks a month has not produced 2,000 outcomes. In the hypothetical below, the same system in the same month looks a third cheaper when it is measured by activity:

HYPOTHETICAL MONTH (ARITHMETIC ONLY, NOT DATA)	COUNT	FULL COST A\$6,000 DIVIDED BY
Tasks attempted	2,000	A\$3.00 per attempt
Tasks completed	1,700	A\$3.53 per completion
Completed and verified correct	1,500	A\$4.00 per verified outcome

The same ratio should be calculated for the current process. A human-only workflow also has a full cost per verified outcome, made of staff time, delay, rework and the tasks that are never completed. Many organisations have never measured it, which is why an agent's price is easier to see than the cost it replaces.

Measure outcomes that are complete and verified. Everything else is activity.

9. Does this need an agent?

An agent can perform almost any step in a clinical workflow. That does not make it the most economic mechanism for every step. Consider an afternoon's requests.

"Mum said someone called about something from her blood test last Tuesday."

This request is genuinely hard. It needs natural-language understanding, identification of the patient being referred to, a check that the caller may receive information about their mother, retrieval of the right result and the right call, and resolution of what "something" and "last Tuesday" mean. Reasoning, and possibly an agent, can add substantial value here. Yet one step inside it is not a reasoning problem: **whether the caller is authorised is a recorded fact, to be checked deterministically, not inferred.**

REQUEST	WHAT IT REALLY IS	MOST ECONOMIC MECHANISM
"Has pathology returned?"	STATE	A query
"Has this result been reviewed?"	STATE WORKFLOW	Workflow status
"Is follow-up overdue?"	TIME + STATE	A rule
"Is the appointment slot still available?"	STATE	A transactional query
"Has consent been recorded?"	STATE	A query
"Send the approved reminder after the defined interval."	WORKFLOW	A scheduled workflow step
"What has changed since the last visit, across letters and results?"	MEANING	A reasoning model, with clinician review
"Mum said someone called about something from her blood test..."	AMBIGUITY STATE	Reasoning or an agent, bounded by deterministic permission checks

An agent can do all of these. That does not mean an agent is the most economic mechanism for all of them.

There is a twist. In a fragmented stack, an agent may genuinely be the cheapest way to answer "has pathology returned?". The question needs no reasoning, but the answer is spread across systems, or locked behind interfaces that only a person, or software acting like one, can operate. **The mechanism a step needs depends on the architecture beneath it.** Paying for reasoning to recover a fact that a coherent system would simply hold is the clearest form of the tax described next.

13 / THE ECONOMICS OF CONNECTED CARE

10. Three taxes

Three concepts help locate where the cost of a workflow comes from. None can yet be measured precisely across healthcare; each is useful for asking where money and time go.

THE REASONING TAX

The recurring cost and operational uncertainty created when probabilistic reasoning is used for a step that could have been reliably deterministic.

- Extra inference and latency
- Variability between runs
- Evaluation and monitoring
- Review and exception handling
- Remediation

It is not “AI is expensive”. Reasoning that changes the outcome is not a tax.

THE FRAGMENTATION TAX

The total cost created because work, information and state are distributed across disconnected systems.

- Duplicate entry and reconciliation
- Switching and searching
- Inconsistent state
- Manual handoffs and delay
- Extra point solutions
- Orchestration to reconstruct context

THE INTEGRATION TAX

The ongoing cost of keeping fragmented systems connected.

- APIs and interface engines
- Connector maintenance
- Version changes and mapping
- Testing and monitoring
- Reconciliation
- Vendor dependencies

The taxes interact, and that tension matters more than any one of them. **An agent can reduce the fragmentation tax, and sometimes the integration tax, while adding reasoning tax.** Whether that trade is worth making depends on the task and on the architecture.

ARCHITECTURE	FRAGMENTATION	INTEGRATION	REASONING	TRANSITION COST
A. Overlay	Can fall: the agent works across systems	Can rise: connectors and screen automation to maintain	Rises where agents handle state and workflow steps	Low
B. Context + action	Can fall: context shared once	Concentrates in the platform	Depends on how much is reasoned	Moderate
C. AI-native core	Can fall furthest, within the platform	Internal interfaces fall; external interoperability remains	Can fall where state is native	High

Table. Directional reasoning, not measurement. Every cell depends on the implementation, and any architecture can be built badly.

The right answer can differ task by task within one organisation. A practice might sensibly use an agent for ambiguous patient messages, workflow for reminders and a query for result status, in the same week.

Agents can move cost from staff time to compute and review. Only architecture can remove it.

14 / THE ECONOMICS OF CONNECTED CARE

11. Where should healthcare spend intelligence?

The debate about agents can make healthcare AI sound like an administrative technology. Some of its most important uses are not administrative at all.

Consider PANDA, a deep-learning model developed in China to find pancreatic cancer on non-contrast CT, a task long regarded as beyond reliable human reading. In multicentre validation across 6,239 patients at 10 centres, it detected lesions with an AUC of 0.986 to 0.996; in the external cohort, its sensitivity for cancers smaller than 2 cm was 92.2%. In a reader study, its sensitivity exceeded the mean of 33 radiologists by 34.1%. Across 20,530 consecutive patients in routine care it found five cancers and 26 other pancreatic lesions that initial standard care had missed. [33] The work is retrospective, has limited data from outside East Asia, and still needs prospective and cost-effectiveness studies.

Newer work looks earlier still. A 2026 Mayo Clinic study reported that a model identified 73% of pancreatic cancers on CT scans taken a median of 475 days before diagnosis, against 38.9% for radiologists reading the same scans. [34] It is a retrospective case-control study, and its specificity of roughly 81% to 88% would produce many false positives in a screening population; the authors call for prospective validation in high-risk groups.

The point is not that clinical AI is good and administrative AI is bad. It is that AI creates value in fundamentally different ways, and the economics differ with them.

CATEGORY	EXAMPLE	WHAT MUST BE TRUE
1. AI compensates for complexity	Navigating fragmented systems	Cheaper than removing the complexity; value may shrink as the architecture improves
2. AI expands clinical capability	Detecting subtle imaging patterns people may miss	Clinical validation in the intended population, regulatory clearance where required, clinician oversight
3. AI interprets human ambiguity	Natural-language requests from patients and clinicians	Grounded in permitted records; permissions checked deterministically; uncertainty surfaced
4. AI assists clinical reasoning	Synthesising history, results and guidance for a decision	Appropriate evidence, governance and human authority over the decision
5. AI should not be used	Checking state, time and defined rules	Deterministic computation is sufficient, cheaper and auditable

Figure 6. Five ways AI can create value, including one where it should not be used. The categories are not a ranking. Category 1 can be worth paying for; its value simply depends on the complexity it compensates for.

Intelligence that compensates for complexity is worth paying for. Intelligence that extends what clinicians can see may be worth far more.

12. Expanding the clinical frontier

Investment in healthcare AI is not only a question of administrative agents. Some of the strongest peer-reviewed evidence of AI extending clinical capability comes from China, where high volumes of routine endoscopy and imaging and large multicentre collaborations have produced trials and validations at scale. The examples below were chosen for the strength of their evidence, not to rank countries.

AREA	STUDY AND DESIGN	FINDING	LIMITS
Gastric endoscopy	ENDOANGEL: single-centre tandem randomised trial, 1,812 patients [35]	Gastric neoplasm miss rate 6.1% with AI first vs 27.3% with routine first	Small absolute numbers; endoscopists not blinded
Upper GI cancer	GRAIDS: 84,424 people across six hospitals of different tiers [36]	Sensitivity 0.942, similar to an expert (0.945) and above trainees (0.722)	Case-control; lower positive predictive value than experts
Gastric cancer on routine CT	GRAPE: 16-centre validation (18,160 people) and 78,593 real-world scans [37]	External AUC 0.927 on non-contrast CT	Retrospective; lower sensitivity for early cancer
Colonoscopy	Randomised trial, 1,058 patients [38]	Adenoma detection 29.1% vs 20.3%	Gain mostly in diminutive adenomas
Diabetic retinopathy	DeepDR Plus: multiethnic development and validation [39]	Could extend mean screening interval from 12 to 31.97 months, with 0.18% delayed detection	Modelled, not a randomised deployment

Combinations, not replacements

The best-evidenced results rarely come from replacing people. A 2026 economic analysis of a real-world diabetic-retinopathy screening programme in China compared eight AI-based strategies across 270 human-AI collaboration scenarios. It found “copilot” human-AI screening the most cost-effective: against the status quo, it was modelled to save US\$1.32 million and gain 14.73 quality-adjusted life years per 100,000 people over a lifetime. [40] It is a model, sensitive to grader and AI specificity and to local costs, and its setting limits transfer to Australia.

Trials point the same way. In Sweden’s MASAI randomised trial, AI-supported mammography screening, with radiologists making every decision, detected more cancers (6.4 vs 5.0 per 1,000) with 44% less screen-reading workload, and interval cancers were non-inferior. [41, 42] The reverse is documented too: across 140 radiologists, inaccurate AI predictions made performance worse. [43]

The evidence favours designed combinations, such as triage, independent double reading and adjudication, over either replacing clinicians or simply showing them a model’s output.

Evidence boundary. These studies concern the cited systems in their own settings. They are not evidence about Echo One, and results from one population may not transfer to another.

13. Intelligence by design

If AI-native means anything, it should not mean AI everywhere. Echo Health AI uses the term to mean **intelligence by design**: build the patient journey and the healthcare operating environment coherently, then use the appropriate computational mechanism at each step.




Figure 7. Intelligence by design. Probabilistic mechanisms propose; deterministic layers hold state and move work; no model output acts without passing typed validation; clinicians hold authority over what matters. The design is combinatory, not maximal.

AGENTS	where dynamic multi-step planning and action create value
REASONING MODELS	where meaning is incomplete, ambiguous, contextual or unstructured
SPECIALISED AI	where machine intelligence can extend clinical capability
DETERMINISTIC SOFTWARE	where state, rules and outcomes are already known
WORKFLOW	where defined processes need to progress reliably
HUMAN JUDGEMENT	where clinical authority, ethics, risk-benefit decisions and accountability belong

Echo One is being built to address two kinds of fragmentation at once: systems (records, booking, forms, messaging, payments, telehealth, results, referrals, billing and portals) and journeys, as patients move between home, general practice, pathology, imaging, pharmacy, allied health, specialists, hospitals and aged care. The patient experiences one journey. Technology too often experiences separate transactions. This is a design direction, not evidence of outcomes.

14. The economics of familiarity

Familiarity has value. Staff know the software. Workflows exist. Integrations exist. Billing keeps running, clinicians stay productive and organisational change is deferred. None of that is laziness. It is an asset, and it has an economic value that belongs in any honest comparison.

Familiarity can also accumulate cost, and one of those costs is rarely counted at all. A decision about the core platform is usually framed as two options: stay or change. It is really three.

COST OF STAYING

- Core platform
- Point solutions
- Integrations
- Agent layers
- Administrative friction
- Duplicated workflow
- Maintenance

COST OF CHANGING

- Migration
- Implementation
- Training
- Workflow redesign
- Validation
- Temporary loss of productivity
- Clinical and operational risk

COST OF WAITING

- More years of accumulated subscriptions
- Continued integration cost
- Continued workflow friction
- Opportunity cost
- Delayed access to a better architecture, if one exists

The cost of waiting is what distinguishes a decision from a default. Every year of delay adds another year of subscriptions, integration and friction to the cost of staying. If a better architecture exists, it also defers the recurring saving that would eventually repay the cost of changing. Choosing to wait is still a choice, and it has a price.

Waiting can also be worth its price. It lets evidence accumulate, lets challengers prove themselves and lets prices fall. The cost of waiting should therefore be weighed against the value of what an organisation learns by waiting, not treated as a penalty for caution.

Two features of the transition evidence belong in this calculation. First, the dip in capacity is planned, not accidental, and recovery takes months rather than weeks. [\[6, 7\]](#) Second, local configuration built up over years, such as decision-support tuning, may not survive a move and has to be rebuilt. [\[8\]](#) Both are real costs of changing, and both are easier to underestimate than subscription fees.

Familiarity is an asset. Unexamined, it is also a liability that compounds quietly.

15. The crossover question

Put the costs over time and the decision takes a recognisable shape. Keeping the current platform and adding to it starts lower. Changing starts higher, because migration, implementation, training and disruption are paid up front. The question is whether, and when, the lines cross.




Figure 8. The crossover question. A conceptual model, not to scale and drawn from no data. There is not yet enough independent comparative evidence to specify the crossover point generically. Every organisation must calculate its own.

Change flattens only if the new architecture actually reduces recurring cost. If it does not, the lines never cross (the dotted line) and change cannot be justified on cost alone. To locate a local crossover, estimate each path over the same three-to-five-year horizon:

1. **Recurring cost of staying:** every subscription including agent seats, integrations, support and maintenance.
2. **Friction:** staff hours on duplicate entry, reconciliation, searching and rework, at a loaded rate.
3. **One-off cost of changing:** migration, validation, training, backfill and months of reduced capacity.
4. **Recurring cost after changing:** from evidence and reference sites, not a vendor estimate.
5. **Risk:** adjust both paths for what could go wrong, including the risk of doing nothing.

16. The burden of proof

Replacing the stack is not automatically better. A replacement platform carries a higher burden of proof than the system it would replace. The incumbent's advantage is inertia. The challenger's obligation is proof. Before it deserves a switch, it must demonstrate:

Safe migration	Data integrity	Uptime	Security
Privacy	Regulatory compliance	Billing continuity	Interoperability
Workflow completeness	Training	Adoption	Clinical safety
Business continuity	Measurable economic value		

That test applies to Echo One. Echo One is being built as a clinical platform with practice management and a patient experience, and its public material is a product preview. This paper is not evidence that it meets any of these tests; that evidence has to come from deployment, measurement and independent review.

What we do not yet know

There is currently **insufficient independent comparative evidence** to show that AI-native unified healthcare platforms have lower total cost of ownership than a legacy record system plus point solutions plus an agentic overlay. Nor is there evidence for the reverse. Vendor claims are not proof, including ours. The market is making competing architectural bets, and the results are not in.

The open questions are specific enough to study:

- What is the full cost per verified outcome of agentic overlays in routine practice, including the review time they create?
- How much of an agent's work is genuine reasoning, and how much is reconstructing state that a coherent system would hold?
- Do AI-native platforms reduce recurring cost after transition, and by how much, in which settings?
- How durable is the value of AI that compensates for complexity, as underlying systems improve?
- What evaluation standard should apply to agents that act inside clinical systems, and who should set it?

The purpose of this paper is therefore to provide a decision framework, not a verdict.

The incumbent's advantage is inertia. The challenger's obligation is proof.

17. A decision framework

Six questions turn the argument into a decision any practice or health service can make with its own numbers.

01 Map each step

Is it state, time, workflow, human ambiguity, clinical perception or clinical judgement? Name it before choosing a tool.

02 Choose the mechanism

Use the least costly mechanism that reliably produces a verified outcome, and reserve reasoning and agents for where they change the result.

03 Measure the full cost per verified outcome

For the current process as well as the proposed one, including review, escalation and remediation.

04 Name the tax

Which tax does this purchase pay: fragmentation, integration or reasoning? Which does it reduce, and which does it add?

05 Price staying, changing and waiting

Over the same horizon, risk-adjusted, with the value of waiting counted as honestly as its cost.

06 Demand proof in proportion to the change

More for replacement than for addition; more for autonomy than for assistance; most for anything that touches a clinical decision.

Conclusion

If we were designing healthcare software today, knowing what we now know about AI, workflows and connected care, would we build the system we are currently paying AI to operate?

For many organisations the honest answer may be no. It may still be right to keep that system for now and add an agent. Both can be true. The cost of standing still is not that incumbents exist, or that agents are priced. It is that the cumulative price of the stack, and the price of waiting, so rarely appear on the same page as the price of change.

Healthcare will need agents, reasoning models, specialised clinical AI, deterministic software, workflow and human judgement, each where it belongs. The durable advantage will not come from using the most intelligence. It will come from spending it where it changes the outcome.

Spend intelligence where it changes the outcome. Build everything else to be known.

21 / THE ECONOMICS OF CONNECTED CARE

References

- [1] **Heidi Health (2026)**. Heidi II puts agents to work on visit admin. Press release, Melbourne, 28 September 2026. Vendor statement. heidihealth.com/en-au/blog/heidi-ii.
- [2] **Heidi Health (2026)**. Heidi II is here. Launch article, 29 September 2026. Vendor statement. heidihealth.com/en-au/blog/heidi-ii-is-here.
- [3] **Heidi Health (2026)**. Why friction is essential for Agents in healthcare. Vendor statement on agent safety. heidihealth.com/en-au/blog/heidi-ii-trust.
- [4] **Heidi Health (2026)**. Pricing (Australia). Screen captures of the live pricing interface held by Echo Health AI; published page checked 29 September 2026. heidihealth.com/en-au/pricing.
- [5] **Huang C, Koppel R, McGreevey JD, Craven CK, Schreiber R (2020)**. Transitions from one electronic health record to another: challenges, pitfalls, and recommendations. *Appl Clin Inform*, 11(5), 742–754. DOI: 10.1055/s-0040-1718535.
- [6] **Rinne ST, Brunner J, Mohr DC, Bearak AC, Anderson E (2023)**. Practices supporting electronic health record transitions: lessons from four US healthcare systems. *J Gen Intern Med*, 38(Suppl 4), 1015–1022. Qualitative study. DOI: 10.1007/s11606-023-08279-0.
- [7] **North F, Pecina JL, Tulledge–Scheitel SM, et al. (2020)**. Is a switch to a different electronic health record associated with a change in patient satisfaction? *J Am Med Inform Assoc*, 27(6), 867–876. DOI: 10.1093/jamia/ocaa026.
- [8] **Wright A, Aaron S, Seger DL, et al. (2018)**. Reduced effectiveness of interruptive drug–drug interaction alerts after conversion to a commercial electronic health record. *J Gen Intern Med*, 33(11), 1868–1876. DOI: 10.1007/s11606-018-4415-9.
- [9] **Anderson E, Moldestad M, Brunner J, et al. (2024)**. User experiences of transitioning from a homegrown electronic health record to a vendor-based product in the Department of Veterans Affairs. *JMIR Form Res*, 8, e46901. DOI: 10.2196/46901.
- [10] **Rigsrevisionen (Danish National Audit Office) (2018)**. Report on Sundhedsplatformen. Government audit, not peer-reviewed. [Read the summary](#).
- [11] **Barnett ML, Mehrotra A, Jena AB (2016)**. Adverse inpatient outcomes during the transition to a new electronic health record system: observational study. *BMJ*, 354, i3835. DOI: 10.1136/bmj.i3835.
- [12] **Holmgren AJ, Apathy NC (2023)**. Trends in US hospital electronic health record vendor market concentration, 2012–2021. *J Gen Intern Med*, 38(7), 1765–1767. DOI: 10.1007/s11606-022-07917-3.
- [13] **Holmgren AJ, Apathy NC, Kanter GP (2025)**. Electronic health record market consolidation and implications for cybersecurity. *Health Aff Sch*, 3(8), qxaf164. Commentary. DOI: 10.1093/haschl/qxaf164.
- [14] **Daniels B, Havard A, Myton R, Lee C, Chidwick K (2022)**. Evaluating the accuracy of data extracted from electronic health records into MedicineInsight. *Int J Popul Data Sci*, 7(1), 1713. The 81% figure is cited from 2016 data. DOI: 10.23889/ijpds.v7i1.1713.
- [15] **Samuelson W, Zeckhauser R (1988)**. Status quo bias in decision making. *J Risk Uncertain*, 1(1), 7–59. DOI: 10.1007/BF00055564.
- [16] **Sinsky C, Colligan L, Li L, et al. (2016)**. Allocation of physician time in ambulatory practice: a time and motion study in 4 specialties. *Ann Intern Med*, 165(11), 753–760. DOI: 10.7326/M16-0961.
- [17] **Arndt BG, Beasley JW, Watkinson MD, et al. (2017)**. Tethered to the EHR: primary care physician workload assessment using EHR event log data and time-motion observations. *Ann Fam Med*, 15(5), 419–426. DOI: 10.1370/afm.2121.
- [18] **Tandem Health (2026)**. Building the AI-native operating system for European clinics. Vendor statement. tandemhealth.ai.
- [19] **Abridge (2026)**. Context-aware clinical intelligence for every clinician at partner systems; company home page. Vendor statements. abridge.com.
- [20] **Ambience Healthcare (2026)**. Company home page. Vendor statement. ambiencehealthcare.com.
- [21] **Innovaccer (2025)**. Why Gravity. Vendor statement, 20 May 2025. innovaccer.com/gravity/why-gravity.
- [22] **Greenway Health (2026)**. Greenway Health launches Novare, the first natively AI-enabled platform built to reinvent the legacy EHR. Press release, 10 March 2026. Vendor statement; efficiency figures from a single vendor pilot. prnewswire.com.
- [23] **Canvas Medical (2025–2026)**. Company home page; Announcing Hyperscribe, 17 March 2025. Vendor statements. canvasmedical.com.
- [24] **Elation Health (2026)**. Elation Health acquires Aster. Announcement, 2 June 2026. Vendor statement. elationhealth.com.

22 / THE ECONOMICS OF CONNECTED CARE

- [25] **mdhub (2026)**. Company home page. Vendor statement. mdhub.ai.
- [26] **Ocean Health (2026)**. Company home page. Vendor statement; early-stage company. meetoocean.ai.
- [27] **Liu Y, Carrero ZI, Jiang X, et al., Kather JN (2026)**. Benchmarking large language model-based agent systems for clinical decision tasks. *npj Digit Med*, 9. Simulated and benchmark tasks; generalist agent frameworks. DOI: [10.1038/s41746-026-02443-6](https://doi.org/10.1038/s41746-026-02443-6).
- [28] **Kapoor S, Stroebel B, Siegel ZS, Nadgir N, Narayanan A (2024)**. AI agents that matter. Preprint, arXiv:2407.01502. Coding benchmark, not clinical. arxiv.org/abs/2407.01502.
- [29] **Jiang Y, Black KC, Geng G, et al. (2025)**. MedAgentBench: a realistic virtual EHR environment to benchmark medical LLM agents. Simulated record environment; early-2025 models. [arXiv:2501.14654](https://arxiv.org/abs/2501.14654).
- [30] **Dratsch T, Chen X, Rezazade Mehrizi M, et al. (2023)**. Automation bias in mammography: the impact of artificial intelligence BI-RADS suggestions on reader performance. *Radiology*, 307(4), e222176. DOI: [10.1148/radiol.222176](https://doi.org/10.1148/radiol.222176).
- [31] **Cottier B, Snodin B, Owen D, Adamczewski T (2025)**. LLM inference prices have fallen rapidly but unequally across tasks. Epoch AI data insight. List prices, not total cost. epoch.ai.
- [32] **Echo Health AI (2026)**. After the Scribe: where the durable value in clinical AI lives. Revised public edition, 25 September 2026. echo-one.health/resources.
- [33] **Cao K, Xia Y, Yao J, et al. (2023)**. Large-scale pancreatic cancer detection via non-contrast CT and deep learning. *Nat Med*, 29(12), 3033–3043. Retrospective development and validation. DOI: [10.1038/s41591-023-02640-w](https://doi.org/10.1038/s41591-023-02640-w).
- [34] **Mukherjee S, Antony A, Patnam NG, et al., Goenka AH (2026)**. Next-generation AI for visually occult pancreatic cancer detection in a low-prevalence setting. *Gut*, online 28 April 2026. Retrospective case-control. DOI: [10.1136/gutjnl-2025-337266](https://doi.org/10.1136/gutjnl-2025-337266).
- [35] **Wu L, Shang R, Sharma P, et al. (2021)**. Effect of a deep learning-based system on the miss rate of gastric neoplasms during upper gastrointestinal endoscopy: a single-centre, tandem, randomised controlled trial. *Lancet Gastroenterol Hepatol*, 6(9), 700–708. DOI: [10.1016/S2468-1253\(21\)00216-8](https://doi.org/10.1016/S2468-1253(21)00216-8).
- [36] **Luo H, Xu G, Li C, et al. (2019)**. Real-time artificial intelligence for detection of upper gastrointestinal cancer by endoscopy: a multicentre, case-control, diagnostic study. *Lancet Oncol*, 20(12), 1645–1654. DOI: [10.1016/S1470-2045\(19\)30637-0](https://doi.org/10.1016/S1470-2045(19)30637-0).
- [37] **Hu C, Xia Y, Zheng Z, et al. (2025)**. AI-based large-scale screening of gastric cancer from noncontrast CT imaging. *Nat Med*, 31(9), 3011–3019. DOI: [10.1038/s41591-025-03785-6](https://doi.org/10.1038/s41591-025-03785-6).
- [38] **Wang P, Berzin TM, Glissen Brown JR, et al. (2019)**. Real-time automatic detection system increases colonoscopic polyp and adenoma detection rates: a prospective randomised controlled study. *Gut*, 68(10), 1813–1819. DOI: [10.1136/gutjnl-2018-317500](https://doi.org/10.1136/gutjnl-2018-317500).
- [39] **Dai L, Sheng B, Chen T, et al. (2024)**. A deep learning system for predicting time to progression of diabetic retinopathy. *Nat Med*, 30(2), 584–594. DOI: [10.1038/s41591-023-02702-z](https://doi.org/10.1038/s41591-023-02702-z).
- [40] **Wang Y, Han X, Hu W, et al., Zhang L (2026)**. Cost-effectiveness analysis of human-AI collaboration in disease screening. *Nature Health*, 1(10). Model-based analysis. DOI: [10.1038/s44360-026-00142-w](https://doi.org/10.1038/s44360-026-00142-w).
- [41] **Hernström V, Josefsson V, Sartor H, et al. (2025)**. Screening performance and characteristics of breast cancer detected in the Mammography Screening with Artificial Intelligence trial (MASAI). *Lancet Digit Health*, 7(3), e175–e183. DOI: [10.1016/S2589-7500\(24\)00267-X](https://doi.org/10.1016/S2589-7500(24)00267-X).
- [42] **Gommers J, Hernström V, Josefsson V, et al., Lång K (2026)**. Interval cancer, sensitivity, and specificity comparing AI-supported mammography screening with standard double reading without AI in the MASAI study. *Lancet*, 407(10527), 505–514. DOI: [10.1016/S0140-6736\(25\)02464-X](https://doi.org/10.1016/S0140-6736(25)02464-X).
- [43] **Yu F, Moehring A, Banerjee O, et al. (2024)**. Heterogeneity and predictors of the effects of AI assistance on radiologists. *Nat Med*, 30(3), 837–849. DOI: [10.1038/s41591-024-02850-w](https://doi.org/10.1038/s41591-024-02850-w).

Publication details

Published by Visualisation Hub Pty Ltd trading as Echo Health AI. Public edition, 29 September 2026. Vendor pages were read on 29 September 2026; products, prices and claims may have changed since. This is a position paper, not a systematic review, financial advice or a product evaluation. Sources support the statements to which they are attached; the frameworks, taxonomy and conceptual figures are Echo Health AI's argument. It contains no patient information.

Companion papers: *After the Scribe, Informed Consent and Shared Benefit and The Digital Front Door—and the Forgotten Back Door*, at echo-one.health/resources. Privacy and security contact: security@echo-health.ai. © 2026 Visualisation Hub Pty Ltd trading as Echo Health AI. All rights reserved.