

PERSPECTIVES ON CONNECTED CARE / 2026

After the Scribe

Where the durable value in clinical AI lives.

A case for connecting information, accountable workflows and patient understanding across the whole journey of care.

Echo Health AI

Visualisation Hub Pty Ltd trading as Echo Health AI

Revised public edition | 25 September 2026

Australian focus. Product direction and illustrative workflows; availability and deployment evidence are confirmed separately.

In this paper

Executive summary	03
1. Scribing is a foundation	04
2. Put investment around the work	05
3. Be precise about risk	06
4. What AI-native should mean	07
5. Integration is an enduring asset	08
6. Medication safety: see the context	09
7. Results: close the loop	10
8. Before the encounter	11
9. Regulation belongs in the design	12
10. The hard problem: holding context	13
11. Four questions before you buy	14
Conclusion: connect the next step	15
References: evidence and limits	16
References: Australian guidance	17

About this edition

This revised edition develops the integration argument of the supplied June 2026 paper. It updates Echo One's platform positioning, replaces the original figures with explanatory diagrams, and removes unsupported market estimates and absolute safety claims. The discussion is a design and evaluation argument, not evidence of Echo One clinical outcomes. Sources are numbered in the text and linked in the reference pages. Companion paper: Informed Consent and Shared Benefit in AI-Assisted Consultations.

01 / THE CONNECTED CARE ARGUMENT

Executive summary

Clinical documentation is an important place for AI to help. A draft note can reduce the work of turning a conversation into a record. It does not, by itself, make the rest of care join up. Someone still needs to reconcile the information, review the result, communicate the plan and take responsibility for the next step.

This paper argues that the durable value of clinical AI lies in those connections. The model contributes language understanding. The surrounding system supplies permitted context, checks what may happen, preserves the record and makes unfinished work visible. A better model can improve part of that chain; it cannot replace the chain.

Three areas make the argument concrete: medication information that must be reconciled across sources; test results that must move from receipt to accountable follow-up; and preparation that helps a patient and clinician begin a consultation with the right context. These are opportunities for careful evaluation, not claims that an AI platform has already reduced harm.

The patient belongs in the same design. A reviewed explanation, an agreed list of next steps and communication that respects the patient's choices can make the outcome of a consultation more useful beyond the consulting room. Consent is a prerequisite; understandable care is part of the intended benefit.

Echo One is being built as a clinical platform with practice management and a patient experience. This edition therefore describes connection both within the platform and across external systems. It does not assume that Echo is always a temporary layer above another practice record.

The useful question is not only what the model can produce. It is whether the whole workflow reaches a reviewed, accountable outcome.

02 / THE CONNECTED CARE ARGUMENT

1. Scribing is a foundation

The attraction of ambient documentation is easy to understand. The clinician wants to attend to the person in the room while producing an accurate record. Software that prepares a useful draft can help with that workload. Published evaluations report benefits in documentation burden and clinician experience, although study designs, products, measures and effects differ. [1, 2]

It would be a mistake to turn those findings into a universal conclusion that scribing is safe, that its benefits are settled in every setting, or that it has no effect on clinical decisions. A note can omit a qualification, attribute a statement to the wrong person or make an uncertain statement sound definite. A convincing draft can influence a busy reviewer even before it is saved.

The practical response is meaningful review against the encounter and relevant source information. Evaluation should measure the quality of the final record and the burden of correction, alongside time. A workflow that produces a quick draft but requires extensive reconstruction has not necessarily saved work.

There is also a limit to what documentation alone can address. The result may arrive after the consultation. The medication list may differ from what the patient is taking. A referral may be prepared without reaching its destination. These problems need connections, ownership and exception handling as well as language understanding.

Scribing remains useful within that larger picture. The argument is to build from it: bring the same discipline of source review and accountable completion to the work that precedes and follows the note.

Evidence boundary: studies of other scribe products do not establish Echo One's performance. Local evaluation must cover the proposed users, setting, model, workflow and release.

03 / THE CONNECTED CARE ARGUMENT

2. Put investment around the work

A healthcare organisation can buy a documentation tool and still leave staff moving information between screens, checking whether messages arrived and remembering who is meant to follow up. A new interface may remove one burden while leaving the handoffs intact.

The investment question should begin with the work. Where does information arrive? Which source is authoritative? What action is required? Who owns it? How does the organisation know that it happened? These questions expose whether a proposed tool closes a gap or creates another place to check.

Integration has a wider meaning here than transferring a field. A result delivered into an inbox is technically connected, but the care process remains incomplete until the right person can review it and record what happens next. Delivery, review and completion are different states. A well-designed workflow makes those distinctions visible.

The same applies to knowledge. An approved patient resource needs an owner and a current version. A practice rule needs scope and an exception path. An external feed needs freshness information and a response to failure. Institutional knowledge becomes useful when it can be retrieved for the right purpose and applied through a governed process.

This is why headline counts of AI features or integrations are poor substitutes for an end-to-end demonstration. Buyers need to see a representative task move through the entire system, including an ambiguous input, a missing source and a failed delivery.

01 Understand the task

Start with the patient or staff need and the present failure point.

02 Connect the permitted context

Identify sources, permissions, currency and missing information.

03 Complete the accountable workflow

Review, assign, act and record the outcome; expose exceptions.

3. Be precise about risk

The existing system has risks: incomplete information, delayed follow-up and unclear responsibility. A proposed AI workflow has risks too: inaccurate content, inappropriate disclosure, misplaced confidence and new failure modes. The harms of the status quo are a reason to investigate improvement, not a reason to waive evidence for the replacement.

Start with the intended use and a bounded population. Identify hazards, estimate their severity and likelihood, and define the controls and evidence required before use. Decide in advance what would cause the workflow to be limited, changed or stopped. Those decisions should be owned by people with clinical and operational responsibility.

Error costs are not symmetrical, but they are not as simple as saying that a false alarm only costs time. Unnecessary escalation can contribute to fatigue, invasive investigation, anxiety or delay for other patients. A missed issue can also have serious consequences. Thresholds and escalation policies need evaluation for the actual context, with both kinds of harm considered.

Observation can be a useful early stage: compare what a system would propose with independently reviewed cases before allowing it to affect the workflow. That requires an approved data environment and appropriate governance. It does not mean copying patient data into an uncontrolled trial.

Once deployed, monitoring should include corrections, unresolved work, failures and unexpected patterns across patient groups. A model or configuration change may alter performance; the safety case has to follow the system that is actually in use.

Risk is managed through intended use, evidence, controls and continuing review. Neither confidence scores nor a human approval button can carry that burden alone.

4. What AI-native should mean

For Echo One, AI-native describes a design approach: the patient journey, information and workflow are designed together, with AI assistance built into that design. It is not a certification, a guarantee of clinical benefit or a claim that every step needs a model.

The model is useful where human meaning is incomplete or difficult to structure. It can help interpret a conversation, prepare a draft or identify a question that needs answering. It cannot confer permission on itself, establish which patient a person means by guessing, or certify that its own output is correct.

Governed context supplies relevant information for an identified purpose. That means knowing where a fact came from, when it was recorded and whether the current user may access it. More context is not automatically better; irrelevant or stale information can confuse the task and expose information unnecessarily.

Deterministic software applies explicit rules to validated inputs: permissions, required fields, appointment constraints, transaction conditions and workflow transitions. It is repeatable, but it is only as correct as its rules and inputs. Clinical content still requires appropriate evidence and meaningful professional review.

01 AI proposes

Interpret uncertain language and prepare a structured draft or question.

02 Context grounds

Retrieve relevant, permitted information with sources and dates.

03 Rules and people govern

Validate the proposal and obtain the required review before consequences.

Architecture principle: model outputs do not act directly. A valid structure is necessary for controlled processing; it is not proof of clinical correctness.

5. Integration is an enduring asset

Models change quickly. The effort of connecting clinical work is different: terminology must be mapped, identities resolved, permissions enforced, external services tested and local workflows understood. Those relationships and processes can remain valuable as the model underneath them changes.

This is the strategic argument for integration. It should not be mistaken for a claim that models no longer matter. Accuracy, latency, language performance and data handling matter greatly. Model choice should serve the workflow while the organisation retains control of its records, policies and responsibilities.

Structured codes illustrate the distinction. A model can propose a clinical concept from a conversation. A maintained terminology and explicit rules can then support a lookup or prepare a related workflow. But a wrong code can be propagated consistently by perfectly deterministic software. The code must be checked against the source and its scope understood.

Nor does a diagnosis alone establish billing eligibility, treatment suitability or the appropriate patient information. Time, service conditions, exclusions and patient circumstances may matter. A mapping is useful only within the conditions for which it has been designed and maintained.

In a connected platform, record ownership must be explicit. Echo One's product direction includes the clinical record and practice operations; external services may remain authoritative for particular information. Each integration needs rules for reconciliation, corrections, duplicate identities, export and downtime. The objective is a coherent record, not competing versions that staff must manually arbitrate.

The lasting asset is the ability to connect trusted information to accountable work, with the provenance and controls intact.

6. Medication safety: see the context

Medication-related harm gives the integration argument a concrete purpose. The Pharmaceutical Society of Australia's 2019 report estimated 250,000 medication-related hospital admissions each year, at an annual hospital cost of about \$1.4 billion. These are historical estimates, not 2026 totals and not a measure of harm attributable solely to disconnected software. [3]

The workflow challenge is familiar: what was prescribed, what was dispensed, what appears on the active list and what the patient reports taking may differ. An allergy entry may be incomplete. A recent change may not have reached every service. A confident summary of one source cannot resolve those differences by itself.

Reconciliation requires sources, dates and clinical judgement. Dispensing data shows supply, not proof of consumption. Absence of a warning is not evidence that a prescription is safe. Dose and treatment decisions require the relevant clinical context and an appropriately governed clinical function.

An illustrative connected workflow would present discrepancies for review, retain the source of each statement, and record the clinician's reconciliation. Approved patient information could then accompany the reviewed plan. That information should not be invented or silently altered by a model, and any adaptation requires appropriate review.

Research on pharmacist-led medication reconciliation supports evaluating this class of intervention, with outcomes and evidence quality varying across studies. It does not prove that an AI implementation achieves the same result. [4] A specific deployment needs its own measures of discrepancy resolution, missed issues, inappropriate escalation and patient outcomes.

Evaluation question: can the reviewer distinguish an authoritative source, a patient report, a model inference and a missing fact before making a decision?

7. Results: close the loop

A test report can exist without the care process being complete. Receipt does not prove review. Review does not prove communication. Sending a message does not prove that the patient understands the plan or that the agreed follow-up has occurred.

Research on test-result follow-up describes failures across several stages and identifies responsibility and communication as recurring issues. Rates vary by study and setting; an international result should not be presented as the prevalence in Australian practice. [5] The useful lesson is to examine the whole process rather than merely the arrival of a report.

Australia's Share by Default requirements commenced on 1 July 2026 for pathology and diagnostic imaging reports authored by, or on behalf of, a pathologist or radiologist, subject to exceptions and extensions. The requirement covers reports, not diagnostic images. [6] Wider availability does not ensure that every result will be present, current or acted upon, and does not establish Echo One's access or conformance.

A connected workflow should make outstanding work visible to the responsible team, with a defined escalation route when it remains unresolved. Changes to a report must reach the review process too; a corrected result cannot be treated as identical to the version previously considered.

01 Received

Record the source, patient match and report version; identify missing or corrected information.

02 Reviewed and assigned

A responsible clinician assesses the result and records the plan and owner.

03 Communicated and followed up

Track the agreed actions, delivery exceptions and evidence of completion.

This diagram is a workflow design, not an autonomous result interpretation or an emergency monitoring service.

8. Before the encounter

The same integration problem appears before the consultation. Patients repeat information already held elsewhere while clinicians start with an incomplete account of the reason for the visit. Preparation can bring those perspectives together without pretending that the system has already assessed the patient.

A useful pre-visit process lets the person explain what matters in their own words, review relevant practical information and understand any proposed use of AI. It should distinguish patient-reported information from verified records. If a model proposes a summary or code, the source and uncertainty should remain available for review.

Preparation and clinical triage have different purposes. A system that recommends urgency, disposition or treatment may influence clinical decisions even when a person sees its output. Its regulatory and safety position must be assessed for that intended use. [7] Calling an interface a front door does not settle that question.

Context assembly also carries risk. The wrong identity, an omitted concern or a misleading summary can affect the encounter. Missing information must not look like a negative finding, and a patient who cannot complete a digital check-in needs another route to care.

The practical aim is a better starting conversation: the clinician knows what the patient wanted to raise, the patient knows what the proposed technology does, and neither has to treat a pre-visit form as a completed clinical assessment.

Preparation is valuable when it helps the right question be asked. It becomes dangerous when a plausible summary hides what remains unknown.

9. Regulation belongs in the design

Australian medical-device regulation turns on the software's intended purpose and the applicable definitions, exclusions and exemptions. TGA guidance distinguishes transcription-related functions from functions that analyse or interpret information for a medical purpose. A digital scribe can fall within regulation depending on what it is intended to do. [7, 8]

Human review does not create a blanket exemption. Neither does calling an output a draft, using fixed mappings after a model, or saying that the software merely surfaces information. Claims made on a website and in marketing material can form part of the intended-purpose evidence. [8]

The design implication is to specify each function clearly. What information does it receive? What does it produce? Who uses it? What decision may it influence? What happens if it is wrong? Those questions should be answered before expanding the use case or promoting a capability.

Clinical governance, privacy, recording consent, health-record obligations and digital-health conformance also require attention. They are related but distinct. A privacy policy does not establish device authorisation; a successful connection does not establish clinical safety; and a certification, if held, has a defined scope.

Echo One's public material is a product preview. This paper makes no claim of TGA approval, ARTG inclusion or blanket exemption, and does not establish that an external health-system connection is live. Deployment discussions need feature-specific intended-use and assurance evidence.

Practical review: agree the proposed use, responsibilities, required authorisations, evidence, acceptance criteria and change process before clinical deployment. This paper is not a legal opinion or a device classification.

10. The hard problem: holding context

A connected system needs enough information to be useful without collecting or retaining everything simply because it might help later. This is a continuing design responsibility. The right answer depends on the purpose, the type of record, the law, the contractual arrangement and the patient's rights.

It is important to separate working material from the clinical record. Audio, transcripts, draft text, signed notes, audit records, messages and backups may need different handling. A promise that all data is transient can be misleading when the service must preserve an authoritative record or evidence of what happened.

For each category, an organisation should be able to explain what is collected, where it is processed, who may access it, why it is retained and how deletion works. Subprocessor and model-provider terms belong in that account, including training use, cross-border processing and downstream copies. Marketing language cannot substitute for those operational facts.

Sovereignty also involves authority and participation, not only a server address. Work involving Aboriginal and Torres Strait Islander health information should be developed with the relevant communities and their governance expectations. A technical design alone cannot establish that custodianship has been respected.

Echo's current public review does not establish a universal hosting location, model-training restriction, clinical retention schedule or independently certified control set. Those matters require approved deployment-specific evidence. The website enquiry retention period also remains unresolved and should not be confused with clinical record retention.

A trustworthy system can explain the information it holds, the authority under which it uses it and the limits of what it knows.

11. Four questions before you buy

These questions are a practical decision lens. They complement clinical governance, procurement, security assessment and implementation planning; they do not replace them.

01 Does the work connect?

Demonstrate the proposed journey across records, people and external services. Identify the authoritative source and the response to missing, stale or conflicting information.

02 Does the patient benefit?

Explain what the person receives: better preparation, a reviewed explanation or a clearer next step. Test understanding and accessibility rather than assuming benefit from clinician time saved.

03 Is the use bounded and evidenced?

Define the population, task, baseline and outcome. Count correction, escalation and failure as part of the workflow. Distinguish studies of other products from evidence about this one.

04 Are responsibilities explicit?

Identify clinical ownership, permissions, consent, intended use, required authorisations and incident handling. Show what happens when the system cannot safely proceed.

An unanswered question is a reason to resolve the gap before the affected use proceeds. It is not something that a feature count, a polished demonstration or a model benchmark can settle.

Measure the complete cost of a verified outcome: implementation, operation, review, rework and support. Record safety and equity effects beside financial measures where a credible monetary value is unavailable. A cheap model call can still sit inside an expensive or ineffective process.

Conclusion: connect the next step

The next phase of clinical AI should be judged by what it helps a care team complete, not only by what it can generate. A useful note is a beginning. Its value increases when it supports a reconciled record, a reviewed plan, an accountable handoff and a patient who understands what happens next.

That is a demanding engineering and clinical task. It requires reliable identity, permissions, terminology, interfaces and operational processes. It also requires restraint: not every uncertainty should be turned into a confident answer, and not every routine step needs a model.

Echo One's direction is to connect practice operations, clinical work and the patient experience within one platform, while supporting the external relationships that care depends on. The architecture must be evaluated in the actual setting. Its ambition is not evidence of safety, availability or measured benefit.

The first step for a buyer is concrete. Choose one representative journey and ask to see how it handles uncertainty, review, responsibility and failure. Then agree what evidence would demonstrate improvement over the current process. Extend the scope only when the evidence and operational readiness support it.

The companion paper, Informed Consent and Shared Benefit in AI-Assisted Consultations, applies the same argument from the patient's perspective: how understanding and choice can begin before a visit, and how useful, reviewed information can continue afterwards.

Connect the information. Make responsibility visible. Prove the outcome.

Explore the product direction and companion publication at echo-one.health/resources. This paper contains no patient records and makes no guarantee of clinical or financial outcomes.

References: evidence and limits

[1] Lukac PJ et al. (2025). Ambient AI Scribes in Clinical Practice: A Randomized Trial. NEJM AI, 2(12). Comparative evidence about ambient documentation; results are specific to the evaluated products and setting. [DOI: 10.1056/Aloa2501000](https://doi.org/10.1056/Aloa2501000).

[2] Olson KD et al. (2025). Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout. JAMA Network Open, 8(10), e2534976. Evidence about clinician-reported burden and experience; not proof of Echo One outcomes. [Read the study](#).

[3] Pharmaceutical Society of Australia (2019). Medicine Safety: Take Care. Historical estimates of medication-related admissions and costs. Figures should not be relabelled as current-year totals or AI-addressable savings. [Read the report](#).

[4] Mekonnen AB, McLachlan AJ, Brien JA (2016). Effectiveness of pharmacist-led medication reconciliation programmes on clinical outcomes at hospital transitions: a systematic review and meta-analysis. BMJ Open, 6, e010003. Evidence about reconciliation programmes, not about an AI platform. [DOI: 10.1136/bmjopen-2015-010003](https://doi.org/10.1136/bmjopen-2015-010003).

[5] Callen J, Georgiou A, Li J, Westbrook JI (2015). The Impact for Patient Outcomes of Failure to Follow Up on Test Results. How Can We Do Better? EJIFCC, 26(1), 38–46. Narrative review of a multi-stage safety problem. [Read the review](#).

This is an editorial whitepaper, not a systematic review. Sources support the particular statements to which they are attached; the proposed architecture and buying questions are Echo Health AI's argument.

References: Australian guidance

[6] Australian Digital Health Agency (2026). Better access to health information. Share by Default requirements for pathology and diagnostic imaging reports, including scope, exceptions and extensions. Page updated 10 August 2026. [Read the guidance.](#)

[7] Therapeutic Goods Administration. Digital scribes. Guidance on when digital-scribe software is regulated as a medical device, with related privacy and safety considerations. [Read the guidance.](#)

[8] Therapeutic Goods Administration. Understanding how we regulate software-based medical devices. Intended purpose, manufacturer responsibilities and the relationship to ARTG requirements, exclusions and exemptions. [Read the guidance.](#)

Publication details

Published by Visualisation Hub Pty Ltd trading as Echo Health AI. Revised public edition, 25 September 2026. Guidance was checked for this edition; requirements and product scope may change. Confirm the current position for the proposed deployment.

This edition preserves the original paper's central argument about integration, medication context, result follow-up, preparation and accountable adoption. It removes unsupported adoption and market-size estimates, unverified quotations and claims of inevitable harm reduction. It also replaces the earlier assumption that Echo is only an overlay with the current connected-platform direction.

The diagrams are explanatory designs. They do not document an approved clinical deployment, regulatory authorisation or live connection. The figures are not patient records or reports of measured Echo One outcomes.

Website: echo-one.health. Privacy and security contact: security@echo-health.ai. Copyright 2026 Echo Health AI.