ARCHORIC | PUBLIC DECISION READINESS

Archoric Decision Notes

Four evidence-backed essays on clinical evidence, implementation, governance and unresolved decision work in healthcare AI.

Confident decisions start before the demo.

ARCHORIC DECISION NOTE 01 | CLINICAL EVIDENCE

Why strong clinical evidence may still leave healthcare AI buyers unable to advance

Clinical evidence can establish that a product deserves serious attention. It does not automatically complete the implementation, governance, security, financial and organisational decision that follows.

4 minute read

Healthcare AI

Clinical evidence can establish that a product deserves serious attention. It does not automatically complete the implementation, governance, security, financial and organisational decision that follows.

A strong study is not the whole decision

Healthcare AI companies are often told to produce more evidence. That advice is reasonable when clinical validity, safety or performance remain weak. But a company can also reach a different problem: the evidence is serious, the product appears credible, and the buying organisation still cannot confidently advance.

This does not mean the evidence failed. It means the evidence answered only part of a larger decision. A clinical leader may understand the reported result while an operations lead still cannot see the new workflow, a security team still needs the data flow, finance still needs a bounded value case, and the internal champion still has to connect all of those views into one defensible recommendation.

The FDA transparency principles make this distinction visible. They call for information about intended use, users, workflow, performance, benefits, risks, limitations, lifecycle monitoring and the needs of different audiences. The important point is not simply that more information is desirable. Different people need different information at different moments in order to make informed decisions.

Evidence availability is not evidence usability

A case study, technical page, security document and regulatory statement may all exist. Yet the buyer may still need to find them across separate routes, interpret unfamiliar language, reconcile different claims and explain the result to colleagues. The evidence is available, but the decision work remains with the buyer.

Archoric calls this the difference between evidence presence and decision usability. A document becomes decision support only when the relevant stakeholder can find it, understand what it establishes, recognise its limits and know what to do next.

This is why adding another general case study can have less value than creating a clear route through evidence that already exists. The problem may not be evidence quantity. It may be sequencing, audience fit, access or translation.

The buyer is evaluating a transition

A performance result describes what the product can do under stated conditions. The buying organisation must also evaluate what it will take to make that result operational. That includes workflow change, integration, training, oversight, monitoring, exception handling and accountability.

The FAIR-AI framework and other healthcare implementation research emphasise multistakeholder review before implementation and continued monitoring after deployment. This reflects a practical reality: a technically credible product can still be difficult to implement responsibly.

The relevant question is therefore not only, "Does it work?" It is also, "Can our organisation introduce, govern and sustain it under our conditions?"

A useful evidence route

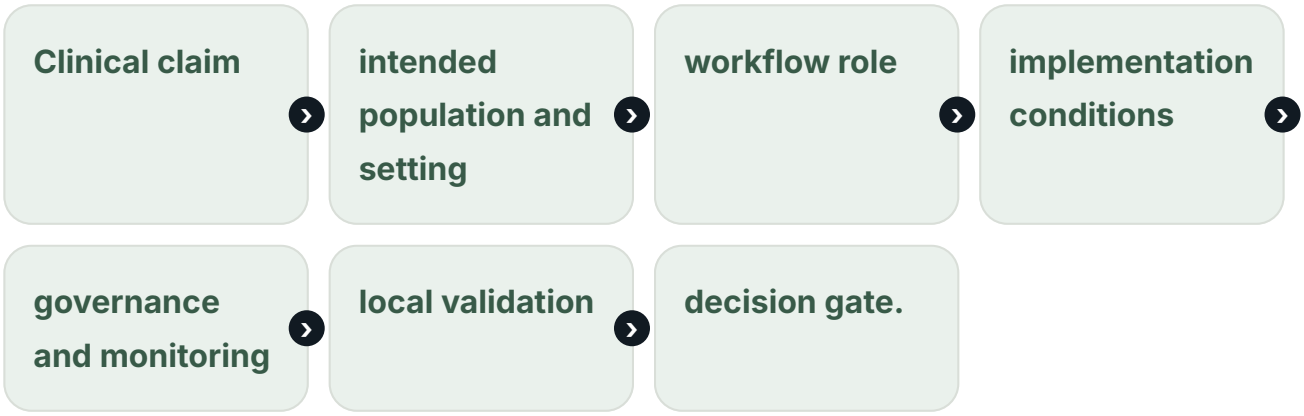

Each step changes what the buyer can responsibly conclude. Skipping a step does not necessarily stop interest. It often transfers the missing work to the champion, who must request another document, schedule another meeting or make an assumption.

What companies can do

First, map every major public claim to the stakeholder decisions it is expected to support. Second, state clearly what the evidence establishes and what it does not. Third, show where public orientation ends and controlled diligence begins. Fourth, build purpose-specific routes for clinical, technical, operational, security and commercial questions. Fifth, equip the internal champion with a concise pack that preserves source links, boundaries and open questions.

This does not guarantee a sale or shorten every evaluation. Budget, product fit, organisational priority and procurement constraints may still determine the outcome. The goal is more disciplined: reduce avoidable decision work and make the remaining uncertainty explicit.

The Archoric perspective

The buyer’s problem may be less about trusting the supplier and more about trusting their own ability to make a complex decision. Strong evidence supports that ability only when it is organised around the work the buying group must complete.

The Proof-to-Decision™ Assessment identifies where public evidence supports or fails that work. The Decision Evidence Navigator™ is designed to turn validated findings into a stakeholder-specific decision environment.

Sources referenced

1. FDA, Transparency for Machine Learning-Enabled Medical Devices: Guiding Principles

2. FAIR-AI, A practical framework for appropriate implementation and review of AI in healthcare

ARCHORIC DECISION NOTE 02 | IMPLEMENTATION

What healthcare AI buyers need to understand about implementation before the demo

The demo can show the product. Before the demo, the buyer is already estimating whether the transition will be workable, governable and worth the organisational effort.

4 minute read

Healthcare AI

The demo can show the product. Before the demo, the buyer is already estimating whether the transition will be workable, governable and worth the organisational effort.

The implementation question begins before implementation

Healthcare AI vendors often reserve implementation details for sales or onboarding. Some information must remain contextual or confidential. But when the public journey says almost nothing about transition, the buyer is asked to enter a meeting without knowing what kind of organisational commitment is being considered.

The buyer may already be asking:

- Which systems are involved?
- Who must participate?
- What changes in the workflow?
- How is local testing performed?
- What happens when the system is unavailable?
- Which responsibilities remain with our organisation?

These are not late-stage technical details. They shape whether the opportunity deserves a next step.

The hidden implementation price

The effective price of a complex health technology includes more than the licence. It can include integration work, internal project time, clinical review, training, governance, monitoring, exception handling, procurement effort and the political cost carried by the internal champion.

A company does not need to publish a universal timeline or fixed workload estimate when those depend on the buyer. It can still make the structure visible. A useful implementation overview distinguishes vendor responsibilities, buyer responsibilities, assumptions, dependencies and the conditions that require local validation.

Five questions public information should begin answering

1. What changes in the current workflow?
2. What systems, data and technical dependencies are involved?
3. Who owns implementation, validation, training and ongoing monitoring?
4. What does a pilot need to prove, and what would stop it?
5. What remains organisation-specific?

These questions help the buyer understand the shape of the transition without pretending that a generic website can complete local diligence.

Implementation evidence should be staged

Visible publicly

the general workflow, implementation phases, prerequisites, typical responsibilities and the route to further review.

Shared after a serious enquiry

technical architecture, integration guidance, resource assumptions, support model and detailed implementation plan.

Confidential under NDA

proprietary technical information, sensitive security details or partner-specific documentation.

Testing for the proposed use case

local acceptance criteria, workflow measures, safety checks, data validation and pilot outcomes.

What companies can do

Create a public implementation orientation page that describes the transition in plain language. Add a responsibility map showing what the vendor provides and what the buyer must validate. Define a pilot as a decision process, not a product trial: name the question, baseline, measures, success criteria, stop conditions and required output.

Finally, connect implementation evidence to the stakeholders who will use it. Clinical leadership needs workflow and safety. IT needs integration and data requirements. Operations needs ownership and exception handling. Finance needs total cost and measurable outcomes. The internal champion needs a route that connects them.

The Archoric perspective

The destination is not enough. The transition has to feel survivable. Archoric examines whether the public evidence gives stakeholders enough orientation to justify deeper review, then designs the evidence routes required for controlled and buyer-specific diligence.

Sources referenced

1. UK Government, A buyer's guide to AI in health and care
2. FAIR-AI, A practical framework for appropriate implementation and review of AI in healthcare
3. Early pipeline framework for assessing vendor AI solutions to support return on investment

ARCHORIC DECISION NOTE 03 | GOVERNANCE

Governance questions healthcare AI vendors should make easier to answer

Governance is not one policy page. It is the practical distribution of decision rights, oversight, monitoring, escalation and responsibility across the product lifecycle.

4 minute read

Healthcare AI

Governance is not one policy page. It is the practical distribution of decision rights, oversight, monitoring, escalation and responsibility across the product lifecycle.

Governance becomes real through decisions

Words such as responsible AI, human oversight and continuous monitoring are common. They matter, but they do not tell a buyer who approves use, who reviews performance, who can override an output, who investigates a concern or who has the authority to stop the system.

NIST describes AI risk management as a way to incorporate trustworthiness considerations into the design, development, use and evaluation of AI systems. WHO guidance similarly places ethics and governance around the full set of actors affected by AI in health. For a buyer, those principles become actionable only when they are translated into decision rights and operating responsibilities.

The governance questions behind a complex purchase

- Who owns approval and oversight?
- Which decisions remain with humans?
- How are model changes reviewed?
- What happens when performance declines?
- How are bias and population fit assessed?
- How are incidents escalated?
- Which responsibilities belong to the vendor, and which remain with the health organisation?

No public page can answer every organisation-specific question. But a vendor can make its own governance model legible and identify where buyer governance must take over.

Governance evidence has levels

Public evidence can explain intended use, human oversight, general monitoring, known limitations and the company’s responsibility model. Controlled evidence can provide governance procedures, validation summaries, change-control detail and incident processes. NDA material may cover sensitive technical or contractual information. Buyer-specific work defines local decision rights, committees, thresholds and escalation.

The value of this structure is not to publish everything. It is to prevent governance from becoming a vague promise followed by a generic contact form.

Governance affects more than compliance

A governance question may begin with compliance but affect clinical safety, workflow, procurement, executive accountability and commercial risk. For example, a model-update process affects technical validation, clinical review, contract notification and the internal champion’s ability to explain whether the original decision still holds.

This is why governance evidence should be mapped across stakeholders rather than owned by one policy page.

What companies can do

Build a governance-to-decision crosswalk. For each governance domain, identify the stakeholder question, the evidence the company can provide, the access level, the responsible owner, the decision it supports and what remains buyer-specific.

Design cross-functional decision meetings for the hardest questions. A model-change meeting, for example, should specify participants, evidence inputs, questions to resolve, required output and the decision gate that follows. This turns governance from a list of commitments into a repeatable decision process.

The Archoric perspective

Archoric does not certify that a company’s governance is legally, clinically or operationally sufficient. It assesses whether the available evidence makes the governance structure understandable, bounded and usable for the stakeholders who must evaluate it.

Sources referenced

1. NIST, AI Risk Management Framework
2. WHO, Ethics and governance of artificial intelligence for health
3. FDA, Transparency for Machine Learning-Enabled Medical Devices: Guiding Principles

ARCHORIC DECISION NOTE 04 | DECISION WORK

What long healthtech sales cycles can and cannot reveal about unresolved decision work

A long evaluation is a signal to investigate, not proof that the website, evidence or sales process caused the delay.

4 minute read

Healthcare AI

A long evaluation is a signal to investigate, not proof that the website, evidence or sales process caused the delay.

Do not turn duration into a diagnosis

Healthtech purchases can take time for legitimate reasons: budget cycles, procurement rules, security review, clinical governance, integration capacity, competing priorities, weak product fit or material risk. A long cycle does not prove that a vendor has a public-evidence problem.

This distinction matters because commercial language often treats delay as friction that should be removed. Some friction is protective. It exists because the decision is consequential and several stakeholders carry different risks.

What duration can reveal

A long or repeatedly stalled evaluation can still reveal where to investigate. Teams can examine whether the same questions return, whether new stakeholders appear late, whether the champion keeps rebuilding material, whether claims require repeated qualification, and whether public information ends before implementation, governance or value questions begin.

LinkedIn and Bain’s recent buyability research describes B2B buying groups as seeking decisions they can defend. That framing is useful here: a credible product and an enthusiastic champion may still be insufficient when the wider group cannot justify the decision within its own responsibilities.

Decision work accumulates

When evidence is scattered, each stakeholder may perform a separate search, request another explanation or reinterpret the same claim. The champion becomes the integration layer. Time is spent not only waiting for approval but constructing the object that can be approved.

This work can appear as more meetings, repeated follow-up, bespoke decks, unanswered security questions, unclear implementation effort or uncertainty about who owns the next action. None of those observations proves causality. Together, they can indicate that the decision environment deserves examination.

A better diagnostic

Separate possible causes into product fit, urgency, budget, procurement timing, implementation capacity, governance uncertainty, stakeholder coordination, evidence availability and evidence usability.

Then record recurring questions by stakeholder and stage. Ask whether an answer exists, where it lives, who owns it, whether it can be shared publicly, and which decision remains blocked. This creates a decision-friction discriminator instead of a generic claim that content will shorten the cycle.

What companies can do

Analyse the last several complex evaluations. Identify where the buyer moved from public orientation into repeated clarification. Map every recurring question to the evidence route that should answer it. Distinguish a missing asset from an existing asset that is difficult to find or difficult to use.

Design the next meeting around a decision, not a status update. Name the participants, evidence inputs, questions, required output and decision gate. If the meeting cannot produce a defined decision, the organisation may still be accumulating information rather than resolving uncertainty.

The Archoric perspective

Archoric does not promise shorter sales cycles. It identifies avoidable decision work that may be contributing to hesitation and builds the evidence environment that helps a company address it. The commercial value lies in making a complex evaluation easier to understand and coordinate, not in guaranteeing the final outcome.

Sources referenced

1. LinkedIn and Bain, The Principles of Buyability
2. Forrester, 2026 State of Business Buying