

AI your institution can defend

Healthcare does not need more AI in the abstract. It needs technology that helps clinicians, patients, caregivers, and institutions make better decisions together – and still makes sense when someone asks what happened.

17%

FEWER REHOSPITALIZATIONS

Across 54 randomized trials of transitional-care interventions, programs that engaged caregivers reduced rehospitalizations 17%. **Programs without caregiver engagement showed no significant reduction.**

Levoy et al., [Medical Care](#), 2022

The interesting part is not the algorithm. It is whether the person going home and the person helping them understood what to do on Tuesday.

Healthcare is full of technically correct information that fails somewhere between the protocol, the clinician, the patient, and the household. That is a human systems problem with clinical consequences.

AI can help. It can also make the distance between those people considerably larger.

60%

of U.S. adults would be uncomfortable if their provider relied on AI to diagnose disease or recommend treatment.

Pew Research Center, 2023

2:1

For every hour physicians spend in direct clinical face time, nearly two additional hours go to EHR and desk work during the clinic day.

AMA / [Annals of Internal Medicine](#)

40%

of asserted medical professional liability cases from 2014–2024 involved communication-related factors.

Candello / [CRICO](#), 2025

HEALTHCARE EXPERIENCE

Lurie Cancer Center

A live oncology research navigator for Lurie Cancer Center Shared Resources, designed so the people doing the work can operate and maintain it without waiting for an AI team.

Simwerx

Offline-capable AI for high-stakes medical training, built to work when there is no cloud available and the answer still matters.

Private deployment

Containerized on-prem or in your VPC for institutions where data boundaries, infrastructure control, and provenance are requirements rather than preferences.

Decide what AI touches

Not every task deserves AI simply because it can be automated. Which decisions stay with a clinician? Which can be supported? Which can be delegated? What evidence does the board need before moving that line?

The first governance decision is not which model to buy. It is what you are willing to let a model do.

In a Black Book Research survey of 182 U.S. hospitals, only 22% reported high confidence they could deliver a complete, auditable AI explanation within 30 days. Black Book Research, 2025

Design for the humans who have to use it

A hospital is not one user. It is physicians, nurses, administrators, researchers, patients, caregivers, risk teams, executives, and families, each with different information, authority, language, incentives, and consequences.

Good healthcare AI recognizes those differences instead of flattening everyone into a persona called user.

That is the anthropology in the architecture.

Make the reasoning inspectable

A clinical answer should have somewhere to point. Protocols and institutional knowledge can be rendered into governed, traceable structures so the clinician can see what grounded an answer, disagree with it, and understand what changed when the answer changes.

That is why we use bounded retrieval, ontologies, and source lineage rather than asking a model to sound sufficiently confident.

A 2026 multicenter validation of a widely deployed sepsis model found substantial performance variation across four health systems. Local context matters. So does knowing when the machine should stop talking. Wong et al., JAMA Network Open, 2026

Tie it to metrics you already report

Readmissions. Triage volume. Discharge comprehension. Time returned to clinicians. Targets are agreed with clinical leads against a named baseline, cohort, and reporting period.

If the work cannot be connected to an outcome your institution already cares about, we should probably not do it.

WHY I DO THIS DIFFERENTLY

I started my career in ruins. Archaeology teaches you very quickly that the official story and the evidence left behind are not always the same thing.

Healthcare has better documentation than most archaeological sites, but the method still holds: look at what people actually do, where information moves, who has authority, what gets worked around, and which artifacts survive the process. Then I became an engineer.

That combination matters here. I can sit with a clinical leader and understand the human system, sit with the technical team and inspect the architecture, and keep following the decision until the two accounts match.

Healthcare AI should not ask people to trust the machine. It should give people enough evidence to exercise their own judgment well. That is a much harder standard. It is also achievable now. **We do not need AGI to do it.**

I keep three advisory relationships at a time. This is a fit when there is a live, consequential decision; an accountable executive; and direct access to clinical, technical, and operational leadership.

It is not a fit for a vendor-selection exercise, a rubber stamp, or advice held at arm's length. Private deployment options, customer-controlled encryption, and BAA documentation are available where applicable. The first call is a conversation, not a pitch.