The Real AI Readiness Question Is Not ‘What Tool?’ Rather It’s ‘What Decision Are We Improving?’

Sit in any life sciences executive boardroom right now and the call for AI innovation is deafening. The standard response from strategic teams is a familiar one: here are the tools, here are the pilots, here is how we help our people operate more efficiently.

The challenge is that we are building upon frameworks which already have systemic flaws. Essentially we are handing out shovels when we need to be architecting a new foundation.

The question most organizations are asking is: What AI tool should we use?

The question we should be asking is: What decision are we trying to improve?

This distinction matters more than most organizations currently recognize, because AI maturity does not start with access to technology. Rather, it starts with knowing which decisions drive quality, speed, cost, and trust. Clarity of purpose and defining the goal is the essential groundwork to intelligent enterprise optimization.

Tool access is not the same as enterprise readiness

Deloitte’s January 2026 State of AI in the Enterprise report, drawing on 3,235 senior leaders across 24 countries including life sciences and healthcare, puts the problem in stark terms: only 34% of companies are using AI to deeply transform their businesses, while 37% are using it at a surface level with little or no change to underlying processes. A healthcare AI leader interviewed in the report captures exactly what I see operationally: without a coherent AI strategy, you get pilot fatigue; we are chasing the next shiny object, pressured to do something with AI without a real plan. Executing a hundred pilots without a clear roadmap leads to poor results and failed value creation.

What we are seeing is the silo effect. We identify individual problem areas without mapping the broader solution and how it impacts the enterprise infrastructure as a whole. We give our teams tools, we teach them to prompt, we open opportunities to pilot new solutions and ask functional teams to identify use cases. However, if each department is solving its own isolated problem without understanding the dependencies across the broader enterprise, we are not building intelligence. We are recreating the current state of fragmented systems with an intelligent layer.

AI transformation does not come from a patchwork of systems. It is true you will see some improvement but the exponential growth and acceleration comes from a holistic architecture which aims at the solution and problem resolution occurs organically with implementation.

Clinical trials do not operate in clean functional boxes.

> A protocol decision affects feasibility → feasibility affects country and site selection → site selection affects activation timelines → activation affects enrollment → enrollment affects forecasting → forecasting affects resource allocation → and resource allocation affects delivery, quality, margin, and leadership confidence.

When AI is implemented as a point solution inside one function, we may improve a task while leaving the real decision problem completely untouched. That is where most organizations are getting stuck. Solving visible problems reactively instead of designing an intelligent operating environment proactively.

The missing layer is decision clarity

In clinical trial delivery, the problem is rarely that we lack data. Enrollment trends, site performance, resource demand, budget actuals, risk signals, deviation patterns, regulatory timelines. The data is usually there and the ambition to leverage the data is present as well. The issue is that it does not connect into a clear decision pathway.

Who needs to know? When? What decision must be made? What dependencies should be visible before that decision is made? What human judgment is required before execution?

Identifying cross functional dependencies, some which will sit completely outside the clinical operation itself, is essential before the build even starts. This is the layer most AI strategies skip entirely.

The FDA's January 2025 draft guidance on the use of AI to support regulatory decision-making makes a point that extends well beyond submissions: AI should be evaluated against the specific decision it supports, the risk of being wrong, and the human accountability required. This principle belongs at the center of every enterprise AI strategy, not just the regulatory function.

You cannot govern a system if you have not defined what decision it is making. We must look to what the end goal is and map the dependencies for clarity.

When you define the objective first, decisions actually land

Our mindset and approach is what usually prevents us from reaching the goals and ambitions we have set. It is easy to focus on the problem and then get bogged down in the details. Instead we need to start focusing on what is the outcome we are looking to achieve and then what are the adjuncts which will impact the final outcome.

Even prior to the age of AI and intelligent computing we were impeded by this nonoperational mental map. I have witnessed this as we have attempted to work through very standard challenges in clinical trials.

A pediatric rare disease program showing real clinical promise but severely delayed enrollment is one place where I saw this mindset directly impact the outcomes. Many of the children we were trying to help were being excluded by their own eligibility criteria. They were too sick to meet the weight thresholds written into the protocol.

There were several valid options to address this challenge: open new sites, expand into new countries, revisit inclusion and exclusion criteria, strengthen advocacy and referral networks. Each solution had advocates, but the failure node had not been anticipated before first patient in. By the time action was required, thinking was dispersed and decision-making was slow. This was not because the team lacked capability, but because there was no framework for determining which intervention would have the most impact, which dependencies mattered most, and where to place energy first.

Research published in Nature using the Trial Pathfinder framework puts hard data behind this problem. When investigators applied a data-driven approach to evaluate common eligibility criteria against real-world patient data, they found that many standard exclusion thresholds had minimal effect on clinical outcomes but were cutting eligible patient pools significantly. Broadening restrictive criteria more than doubled the eligible population on average. The science was not driving those exclusions. It was outdated habitual thinking and assumption that were.

I also led a team where we were tasked with activating 30 sites by a hard deadline. Here we implemented the use of a basic LLM and we defined the goal first. Through defining the goal we worked backward: which sites had the fewest outstanding requirements, which countries had the most predictable regulatory behavior based on historical patterns, where could we move fastest without sacrificing quality. We identified potential failure nodes and strategically addressed them, thus helping us to achieve the goal. The difference was not the tool. It was that we had defined the decision framework before we touched the technology.

This contrast is the point and it demonstrates why our partnership with AI tools and systems is so important. In both scenarios the data existed. What was different was decision clarity upstream.

This same principle applies at the most foundational level of trial design.

Protocol optimization is still the black box

If we want to build true systemic intelligence, we have to start at the foundation: the protocol.

A 2025 perspective published in npj Digital Medicine examining AI, LLMs, adaptive trial designs, and digital twins in clinical research identifies a consistent finding across the literature: the persistent challenges in cost, enrollment, and generalizability are not being solved by individual tools applied to isolated problems. The infrastructure connecting those tools, and the decision logic that runs through it, is what is missing.

We must define the objective goal as our starting point and then determine the dependencies to reach this goal. This goes beyond digitizing the protocol for operational efficiency. We are actually looking at optimizing the protocol and introducing advanced analysis techniques to support our defined outcome. The achievement of this extends beyond a single AI solution but rather incorporates an intelligent infrastructure to optimize across several functional dependencies.

Optimization asks core questions such as: Can this protocol actually enroll? Which eligibility criteria create avoidable barriers? Where will site burden exceed site capacity? Which endpoints introduce operational complexity that will surface six months into execution? What risks should be resolved before the trial ever launches?

Imagine applying knowledge graph architecture to map every dependency within a protocol: site selection, regulatory considerations, patient burden, cost-effectiveness, country feasibility,

endpoint risk. This moves beyond fixing a workflow but develops an ecosystem which supports endpoint achievement and effective protocol execution. You are designing a system that anticipates failure nodes before the trial launches. This shifts the industry from reactive rescue to proactive design, and in clinical development the shift is worth years.

The enterprise is not a collection of problems

We have to view the enterprise as a single fluid body. You do not treat the circulatory system in isolation and expect the kidneys to benefit.

A forecasting tool alone will not fix resource planning if staffing decisions stay disconnected from operational demand. A finance dashboard alone will not protect margin if project managers are the ones seeing scope creep first but only finance is empowered to act. A risk tool alone will not improve delivery if signals are captured but never translated into timely decisions.

Deloitte's 2026 research is direct on this point: the most successful organizations are not those optimizing what already exists. They are rethinking their organizations from the ground up, building without legacy constraints rather than digitizing old processes. The gap between availability and adoption is now the primary barrier to value, and closing it requires designing for deployment from the outset, not treating scale as an afterthought.

Intelligent enterprise optimization represents an opportunity larger than tool deployment. It is the design of an intelligent operating environment where systems share context, dependencies are visible, and human experts are supported at the point of decision. Knowledge graphs, machine learning, NLP, LLMs, and governed agentic frameworks become genuinely powerful not because they are impressive individually, but because they can map relationships across complex interdependent systems. In clinical development, those relationships are everything. One decision is rarely isolated. Every decision creates downstream consequence.

The leadership question needs to change

The most common question I hear from senior leaders is: "How will this be different from past attempts?"

It is a fair question. But the answer should not start with a feature comparison. It should start with approach.

Are we solving a pain point or redesigning the operating model that keeps producing it? Have we defined the ideal future state before building the problem list? Have we mapped dependencies from enterprise strategy down to daily execution?

Start with the goal. Define the decision. Map the dependencies. Identify the failure nodes. Build the infrastructure. This creates the solid foundation for substantial change and the capabilities to meet operating expectations over the next 10+ years.

If you are a VP of Clinical Operations evaluating your next vendor pitch, come to that conversation with your ideal operational end state already mapped. Define the decisions that matter most to your enterprise and push the vendor to meet that vision. Do not let them sell you their product roadmap when what you need is a systems partner.

The organizations that lead in AI-enabled drug development will not be the ones with the largest tool catalogs. They will be the ones who understand their business deeply enough to redesign how decisions get made, and bold enough to build the infrastructure to connect them.

In clinical development, this requires leaders who can bridge deep operational domain expertise, AI strategy, governance, and enterprise systems thinking. This is not a niche technical skill. It is

the most important leadership capability in this industry right now and it comes from critical experience of building domain expertise.

The future will not be won by asking what tool to buy. It will be won by leaders willing to ask: what decision are we improving, and what kind of enterprise do we need to become in order to improve it?

References (first comment) Deloitte AI Institute, *State of AI in the Enterprise: The Untapped Edge*, January 2026 | BCG, *AI at Work 2025: Momentum Builds, But Gaps Remain*, June 2025 | Xu et al., *Evaluating Eligibility Criteria of Oncology Trials Using Real-World Data and AI*, Nature, 2021 | Giuffrè & Shung et al., *AI and Innovation in Clinical Trials*, npj Digital Medicine, November 2025 | FDA, *Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products*, Draft Guidance, January 2025