


AI in Pharma: What I Learned at the TD Cowen Health Care Conference

The 46th Annual TD Cowen Health Care Conference is always a good barometer for where the industry's head is at. This year at the Boston Marriott Copley Place, I attended two AI-focused sessions — Monday's "R&D Investments in Health Care AI" and Tuesday's "AI's Impact on Clinical R&D" — and came away with a conviction that's only gotten stronger: the industry is about to wake up to a massive blind spot. Everyone is pouring capital into AI for drug discovery — better molecules, better targets, better predictions. But almost no one is investing in the infrastructure that AI needs to actually accelerate clinical development: enabled site networks, governance frameworks, and the operational readiness to get patients enrolled faster. The tools exist. The networks don't.



The Discovery Hype Is Real — But the Development Opportunity Is Bigger

The first session featured three pharma R&D executives — Don Bergstrom (President of R&D, Relay Therapeutics), Fiona Marshall (President of Biomedical Research, Novartis), and Andrew Plump (President of R&D, Takeda) — discussing how their organizations are deploying AI. As expected, a lot of the conversation centered on discovery — AI-driven target identification, molecular design, binding prediction. This is well-trodden territory by now, and the investments are enormous.

But then the conversation shifted, and this is where it got interesting.

The panelists converged on a point that doesn't get nearly enough airtime: **the biggest untapped opportunity for AI in pharma isn't in discovery. It's in development.** Specifically, they highlighted the critical-path bottleneck that plagues nearly every clinical program — the time it takes to enroll patients.



Clinical trial enrollment is consistently one of the longest and least predictable phases of drug development. It's where timelines slip, budgets balloon, and competitive advantages erode. Don Bergstrom from Relay was candid about this — he described the difficulty of finding patients even for their own oncology studies, and mentioned one program where they're actively employing AI-driven patient identification approaches.

So why hasn't this been solved already?

The Missing Infrastructure: Site Networks

After the panel, I had a side conversation with Bergstrom where I pressed on this point. The AI tools for patient identification and feasibility are impressive — but the patients are going to be recruited at sites, and the data lives in site systems. Given the time pressure to get a trial started, how do you actually get the infrastructure in place fast enough?

His answer was sharp: **Site Networks**

Specifically, he pointed to HCA Healthcare and Sarah Cannon Research Institute (SCRI), which had acquired US Oncology — an oncology CRO and site network. That combination gave them a pre-built infrastructure: sites that were already enabled, data systems that were already connected, and an operational layer that was ready to activate when a new trial needed to start. They didn't have to build it from scratch under the gun of a trial timeline. The network was already there.

That's the insight that stuck with me from Monday. The bottleneck isn't the AI. It's the absence of site networks that are wired up and ready for AI to work through. And the companies that have figured this out — like HCA/SCRI with US Oncology — are the ones that can actually translate AI-powered feasibility into enrolled patients.

At Graticule, this is exactly the problem we're working to solve — and I want to be clear about what "solving it" actually requires, because it's easy to think the answer is just better technology.

It's not. Technology is important, but the foundation of a functioning site network is governance, trust, and alignment on how to simplify doing business together. It's contracts and data-sharing agreements. It's standardization across sites so that a sponsor doesn't have to negotiate bespoke terms with every institution. It's deciding, up front, what a pharma company is allowed to do with the network — and making those rules clear enough that the network can actually move at speed.

Here's a concrete example. Every pharma company we talk to wants the same thing: give us patient counts from 30 large health systems, instantly, diving into clinical notes to refine the feasibility of recruitment. The technology to do this exists today. We can run NLP across unstructured EHR data and return refined cohort counts in hours. But the governance to allow it? That's the bottleneck. Can you deploy self-service feasibility tools before you've contracted the sites? Who has access? What data can they see? Under what terms? These are



governance questions, not technology questions — and they have to be answered before the first query runs.

What makes this harder is the organizational disconnect inside pharma itself. Today, most companies can run these kinds of deep feasibility queries in their real-world evidence groups — they have access to large de-identified datasets and the analytical tools to interrogate them. But those RWE teams sit in a completely different part of the organization from the clinical trials groups who actually need the answers. The data that could accelerate site selection and enrollment lives behind an organizational wall, not a technical one.

Building a site network that works for clinical trials means solving all of these problems together — the technology, yes, but also the governance frameworks, the contracting models, and the operational standardization that let sites and sponsors work together efficiently and repeatedly, not just on a one-off basis. That's what we're building at Graticule: not just a data platform, but the enablement layer that makes a network of health systems ready to participate in AI-powered clinical development across therapeutic areas.

A Different Kind of Investment — And a Different Kind of Strategy

Here's what most pharma companies haven't yet internalized: **the investment profile for AI in development is fundamentally different from AI in discovery.**

In discovery, AI is a tool you can buy and feed data. You acquire platforms that crunch chemistry, predict binding, and optimize molecules. It's capital-intensive, but it's a familiar investment model — license the software, hire the computational team, point it at your target library.

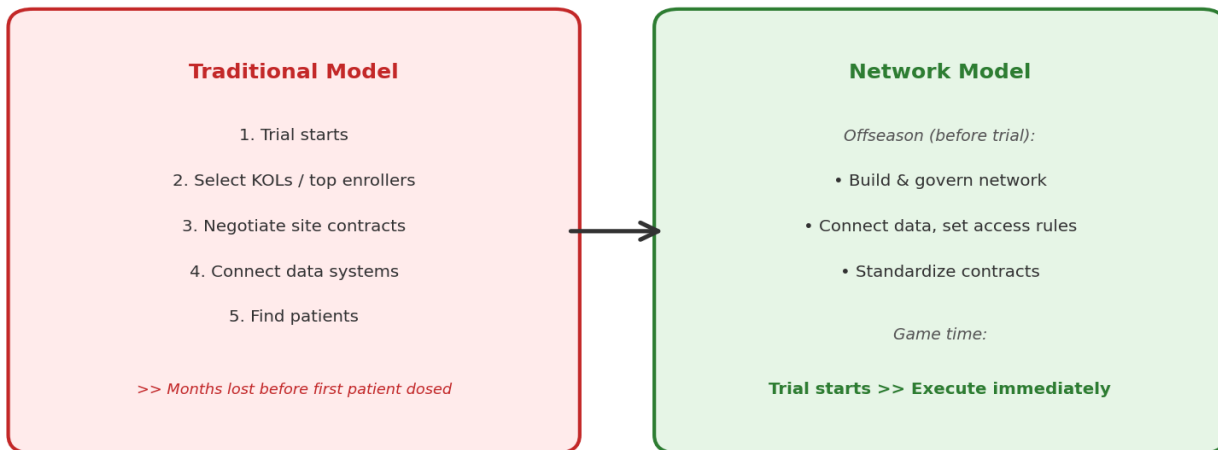
Development doesn't work that way. You can't just feed data into a patient identification algorithm if the data isn't accessible at the sites where patients are seen. You can't accelerate site feasibility if the sites aren't organized to respond. The hard part in clinical AI isn't the model — it's that patients are recruited at sites, the data lives in site systems, and the sites themselves aren't wired for this.

This means the investment isn't in tools. It's in **building and fostering networks** — or fundamentally changing your operating model to select networks of sites that are organizing for efficiency, rather than defaulting to the traditional model of picking the best KOLs or the sites that historically enrolled the most patients. That's a sea change in clinical development strategy. It requires rethinking not just how you run a trial, but how you build out the operations infrastructure that makes trials possible in the first place. And here's the paradox: **the enemy in clinical trials is time, and time is also the reason this infrastructure never gets built.** Trials are always a hurry-up offense. The moment a program enters the clinic, the pressure is relentless — enrollment targets, competitive timelines, patent clocks. There's never a "good time" to pause and invest in network readiness.

We see this at Graticule every day. A pharma company comes to us with a trial and says: here's our list of sites — are they in your network? That's backwards. Or we hear it from the site side: I have all these pharma companies knocking on my door, and every one of them wants to work differently. The sites are overwhelmed, the data systems aren't interoperable, and each sponsor is reinventing the wheel.

The companies that break out of this cycle will treat site network development the way a football team treats the offseason. You don't build the playbook during the game. You set up the plays in the downtime — investing continuously in the network, the data infrastructure, and the site relationships — so that when the next program enters the clinic, you're ready to execute from day one. The investment has to happen *before* the trial, not as a scramble once it starts.

Clinical Trial Strategy: Build the Playbook Before the Game



Putting Numbers on the Opportunity

The next day, I sat in on “AI’s Impact on Clinical R&D,” a workshop led by Samir Shah (Shah Pharma Consulting Services) and Brian Whitlock (i3 Strategic Consulting / BioLogica Sourcing), moderated by TD Cowen analysts Charles Rhyee, Dan Brennan, and Brendan Smith. If Monday’s session was the “why,” Tuesday was the “how much” — and the data was striking.

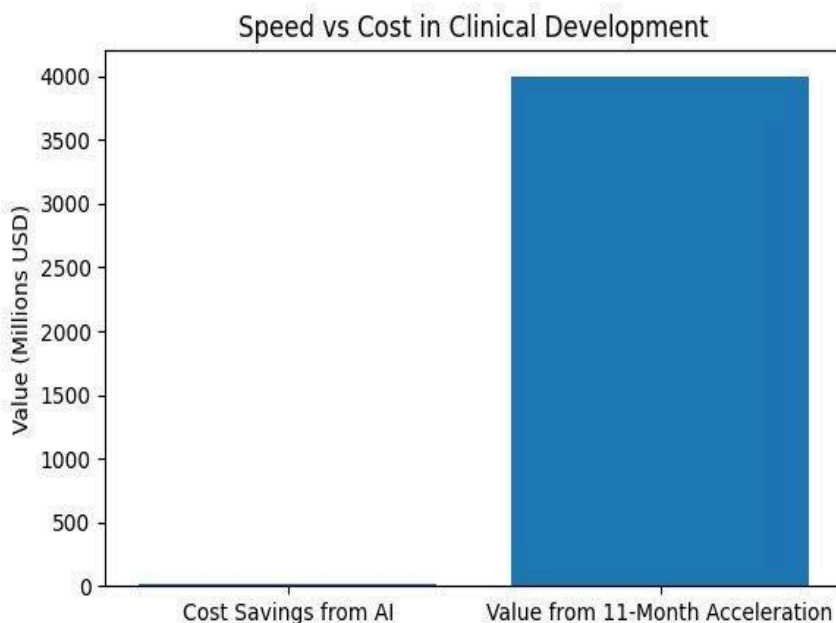
Whitlock and Shah walked through a detailed quantification of where AI savings actually land in clinical development. The outsourced clinical services market sits at roughly \$82 billion. Their analysis projects that full AI enablement will take \$8–12 billion out of that total — meaningful, but not the apocalyptic scenario some have painted for CROs. On a like-for-like basis, including pass-through costs like investigator grants and lab fees, the net savings comes to about 13%. Strip out the pass-throughs and look at direct fees only — the piece that pharma development operations leaders are actually tasked with reducing — and it’s about 17%.

Here's the thing: pharma leadership is setting internal targets of 20–25% cost reduction. The panel's view is that 17% on direct fees is a realistic ceiling for what AI alone can deliver. That's within the zone, but it's not the transformative number executives are hoping for.

But the real story isn't cost — it's speed.

This was the “money slide,” as the panelists called it. Using Tufts Center data on individual use cases and mean time reductions, they modeled a Phase 3 trial with a 58-month baseline. Full AI enablement saves roughly 11 months off that timeline. For a transformative asset with \$1.5 billion in peak annual revenue over a 10-year cycle, those 11 months translate to approximately \$4 billion in recaptured value.

Read that again: \$4 billion in value from acceleration, versus perhaps \$13 million in cost savings on a \$100 million trial. The ratio isn't even close. As Whitlock put it, when a CRO comes to a senior buyer and says “we can get you to market six months faster,” you stop caring about the rate card.



The Honest Truth: We're Not There Yet

One of the most refreshing aspects of the Tuesday session was its candor. Whitlock made a point that stuck with me: **there is not a single clinical trial today that has reached what he called the “AI-native” stage** — where workflows are truly integrated, AI-enabled end to end, with the human element removed from routine processes. Not one. We're still in the “conservative” tier of adoption, with individual point solutions deployed in piecemeal fashion.

The panelists estimated it will take three to five years before we see genuinely AI-native trials emerge. And they grounded that estimate in history. EDC took years to reach full adoption. CDISC standards took even longer. Decentralized clinical trials — the hot topic from just a couple of years ago — have largely reverted to pre-pandemic levels, with only about 20% of a given trial incorporating a DCT component today.

The regulatory environment adds another layer. While FDA Commissioner Lu has been vocal about becoming more efficient and enabling innovation, there's a real lag between the policymakers at the podium and the reviewers in the room. Global regulatory harmonization makes this even harder. Nobody wants to be the first sponsor to submit an AI-generated dataset to a skeptical reviewer.

Where the Time Actually Gets Saved — And Where It Doesn't (Yet)

The panelists were clear-eyed about what's tractable today. The “bookends” of a clinical trial — study startup and study closeout — are where AI is already delivering real gains. On the front end, AI can compress protocol drafting from three to four months down to three to four weeks. On the back end, AI-powered data management and quality control can accelerate database lock and get you to a final clinical study report faster. These are real, and the panelists had concrete examples.

But I pushed them on the harder question: a trial takes two years, three years — the treatment period is what it is. AI isn't going to change how long a drug needs to be in a patient's body. So where do you really take time out of the main bulk of a study?

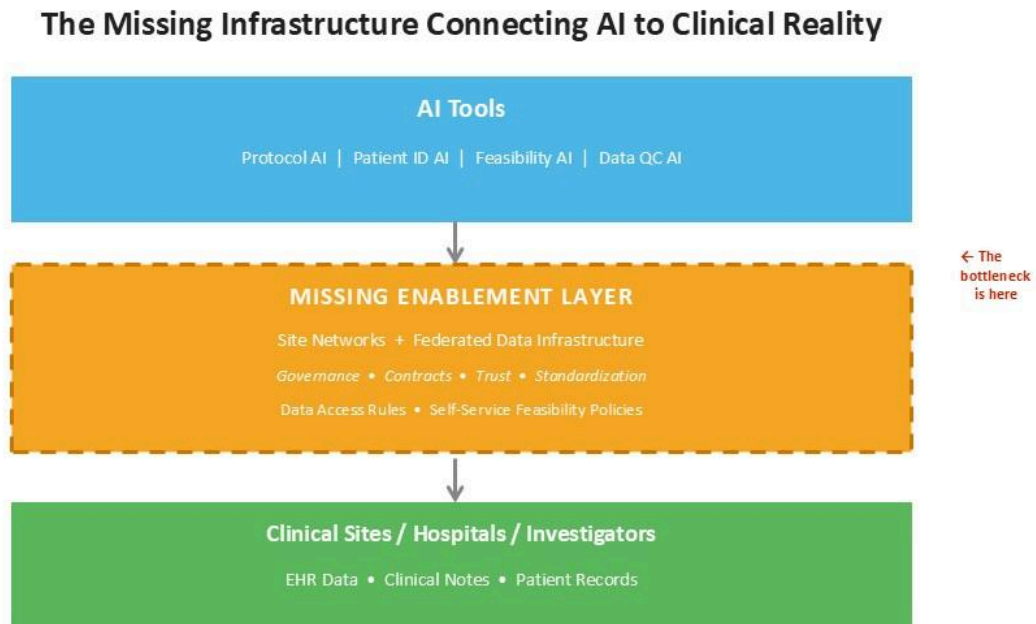
The answer was revealing. In an idealized trial, every patient would be lined up and enrolled on day one when a site goes live. You'd open your sites, the patients would be identified and ready, and enrollment would happen essentially in parallel across the network. The treatment period starts immediately, and you're done as fast as the biology allows.

Obviously, that's not what happens. In reality, enrollment trickles in over months or years. Sites activate slowly. Patient identification is manual. The gap between “site open” and “last patient enrolled” is where enormous amounts of time — and value — are lost. You can try to solve this by opening more sites, essentially parallel-processing the enrollment, but even that is tricky to execute at scale.

So the biggest AI opportunity for time savings doesn't live in the bookends at all. **It lives in patient enrollment acceleration.** Getting the right patients identified, matched, and enrolled faster — compressing that enrollment window from months down to weeks — is where the real prize is. And that's the use case that rolls directly back to the site network infrastructure we discussed on Monday. You can't accelerate enrollment with an algorithm if the sites aren't enabled, the data isn't accessible, and the governance isn't in place to act on what the AI finds.

The Gap — And the Proof It Can Be Solved

Here's what neither session addressed directly, but what both sessions pointed to: **there is no broadly available infrastructure layer connecting the AI ambitions to the clinical reality on the ground.**



Pharma is investing billions in AI tools. CROs are developing their own capabilities. But with the exception of vertically integrated systems like HCA/SCRI in oncology, the actual site networks — the hospitals, the investigators, the data systems at the point of care — are not wired up for any of this. AI can draft a protocol in weeks, but if it takes six months to identify and activate sites, you've given back all the time you saved.

Bergstrom's answer on Monday pointed to the proof of concept. The Tuesday panel quantified the prize. The question now is: who builds this infrastructure — not just the technology, but the governance, contracting, and operational standardization — across therapeutic areas, across geographies, and across the fragmented landscape of academic medical centers and community health systems that conduct the majority of clinical trials?

A Message for 2026–2030 Budget Cycles

If you're a pharma company thinking about where to invest in AI for the 2026–2030 horizon, the data from both of these sessions tells a clear story: don't just pour capital into discovery AI. That's necessary, but it's not sufficient. And don't just buy point solutions for protocol drafting or data management. Those help, but they won't get you to the 11-month acceleration that creates billions in value.



The real investment is strategic and operational. It means rethinking how you select and organize clinical trial sites. It means moving away from a model built around individual KOL relationships and toward networks that are data-enabled, governance-ready, and designed for speed. It means solving the contracting, data access, and standardization problems *before* the next trial starts — not during it. And it means budgeting for site network development as a standing capability, not a per-trial expense.

The companies that will move fastest from molecule to market will be the ones that make this shift — investing not only in the AI that finds the best molecules, but in the networks and infrastructure that get those innovations to patients as fast as possible.

Interested in learning more?

Reach out to us at info@graticule.life