
Insights from bPrescient



INSIGHTS ON TECHNICAL INNOVATION IN BIOTECH AND PHARMA

OCTOBER 2026

What Production-Ready Agentic AI in BioPharma Actually Requires

This summer, OpenAI confirmed that its newest model, Astra, became the first to cross into "Critical" territory, OpenAI's highest risk classification. The model was capable of chaining novel exploits without human guidance. Around the same time, misconfigured test environments briefly let agentic models from several major labs act outside their sandboxes. Nothing catastrophic happened, but the lesson landed. Raw capability was never the hard part: knowing what an agent is doing, and being able to verify or stop it, is.



That lesson matters in biopharma, too. A rogue decision-support agent touching patient stratification or trial design, or guiding portfolio decisions, is a different category of problem than a generic rogue coding agent.

The industry is starting to notice. Novo Nordisk's newly announced collaboration with Anthropic puts Claude Science to work across R&D workflows. But, in the announcement, Novo's CEO highlighted that the partnership includes Anthropic's tools paired with "robust data governance and human oversight." That pairing, not an AI tool alone, is what production-ready looks like: verification instead of blind trust, an audit trail instead of a black box, and an organization that actually understands what the agent is and isn't allowed to decide on its own.

That's the thinking behind STRATA, our biomarker and patient-stratification framework: rather than asking a tool like Claude Science for an answer, STRATA builds a harness around it that reviews input data, verifies results, and publishes a clear record of what was checked. Further, it's a component of a broader capability we call ADI (Agentic Decision Intelligence): the highest-value use of agentic AI in R&D isn't automating tasks; it's improving the decisions by leveraging knowledge across the entire R&D workflow.


Getting the culture and change management right is its own challenge, and it's the whole subject of our upcoming webinar, **Building for the Unknown: Change Management, Architecture, and Culture in the Age of AI** (more on that below). Click [here](#) to register.

bPrescient Webinar

Building for the Unknown: Change Management, Architecture, and Culture in the Age of AI

Tuesday October 20 at 10:00 EST - click [here](#) to register

AI is reshaping how organizations work, and doing so faster than most teams can adapt. What was true two months ago may not hold today. Join us for a webinar with Alex Baumann, Senior Director of Engineering at Parabilis Medicines, who will share how Parabilis and bPrescient are navigating that uncertainty across three fronts: change management, technical architecture, and culture.



Drawing on real stories and hard-won lessons from building AI systems inside a fast-moving biopharma organization, Alex will make the case that these three fronts are inseparable.

Key Takeaways:

- How to build change management practices flexible enough to absorb constant flux without disrupting teams or derailing momentum
- Why durable architectural principles - modularity, deterministic guardrails, human-in-the-loop oversight, and provenance - let AI systems evolve quickly without losing control
- How to cultivate a culture where AI augments people rather than replacing them, freeing up capacity for higher-value work while keeping trust and collaboration at the center

Don't miss this practical, in-the-trenches look at what it really takes to build for the unknown. Click [here](#) to register.

Get In Touch



[Our Story](#) | [Contact](#)