
Insights from bPrescient




INSIGHTS ON TECHNICAL INNOVATION IN BIOTECH AND PHARMA

JULY 2026

From Pilot to Production: Why Agentic AI Stalls at Scale

In the past year, many of bPrescient's clients have developed agentic solutions and seen the impressive promise of this new technology. Their question now is why so few of their successful pilots and proofs of concept ever make it into dependable, day-to-day production.



It's not developing the agents themselves that stalls deployment; it is everything around them: governance, validation, integration, change management, and making the leap from a controlled demo to a large and varied user base.

As an illustration, one of our clients has developed several agents to aid in analysis of DNA sequence data and other molecular data for translational research. The agents were very useful in the hands of the developers, but the effort faltered when they tried to make them more broadly accessible across multiple teams.

We see three obstacles that occur repeatedly:

- **Governance.** A pilot run by a small, expert team carries little risk. A production agent that can pull data, invoke tools, and take actions across multiple workflows is a different matter, especially in regulated environments. Governance like this is difficult to add after the fact; it has to be designed in from the beginning.
- **Validation.** A demo that performs well on clean, curated inputs tells you little about behavior at scale. Production-grade agents require defined performance criteria, drift monitoring, clear subject matter expert review points, audit logs, and the provenance to reconstruct any decision: what data went in, which model version processed it, and who reviewed the output.
- **Generalizability.** Pilots succeed under narrow conditions with a few well-understood data sources and a handful of expert users. Real users bring messy data, edge cases, and expectations the original design never anticipated. An agent that meets the needs of a few power users can frustrate other users if its scope, guardrails, and failure modes were never defined for that broader audience.

None of these obstacles is a reason to slow down. The solution is to build differently. Organizations that are pulling ahead treat agent deployment as an operational discipline rather than a technology demo: they map data and ownership before writing agent code, embed evaluation and human oversight from the outset, and design reusable, governed architectures that each new agent can operate safely in.

This is bPrescient's focus. *We help research and clinical teams move from compelling pilots to validated, governed, production-ready agentic systems by building the evaluation frameworks, oversight, and standardized architectures that make scaling both safe and sustainable.* If your agents are succeeding in the hands of a few experts but are stalling on the way to broader production, email us at info@bprescient.com.

[Get In Touch](#)