



Reuters Pharma USA Roundtable Learnings: Real Barriers and Real Solutions to AI Adoption in Life Sciences

Christopher Shipp | Founding Partner, Humigent | March 2026


Applied AI | Glass Box Intelligence | Commercial Strategy
in Life Sciences





Key Findings

Based on a facilitated executive roundtable with nearly two dozen pharma leaders, private client conversations, and observations from the full two-day conference, this paper argues that the primary barriers limiting AI adoption in pharma are not wholly technical. They are human and organizational. The technology is ready. The question is whether the organizations, teams, and individuals expected to act on it are ready too.





Background

Reuters Pharma USA 2026 brought more than 1,500 leaders to the Pennsylvania Convention Center in Philadelphia in mid-March. Commercial, analytics, market access, medical affairs, field operations, and digital transformation functions were all well represented. It is one of the industry's most reliable check-ins on where pharma stands, and what the real conversations look like beyond the press releases.

Humigent joined as an exhibitor and as the host of an executive roundtable titled *"From Hallucinations to Reliable, Auditable AI: The Case for Glass Box Intelligence."* The roundtable drew 25 registered leaders from AstraZeneca, BMS, Eli Lilly, Johnson & Johnson, Merck, Novartis, Pfizer, Shionogi, UCB, Argenx, Gilead, and Memorial Sloan Kettering Cancer Center.

The hypothesis going in was straightforward. In a regulated industry where decisions carry both clinical and commercial consequences, AI must meet two requirements: reliability (accuracy, consistency, no hallucinations) and auditability (the ability to trace and prove how any output was generated). The expectation was that these two concerns would dominate the room.

They did not.

An Unanticipated Insight

Under Chatham House rules, the conversation became candid quickly. What emerged from the group, consistently and across functional roles, was not primarily a story about model accuracy or audit trails. It was a story about people and organizations.

The pattern was strikingly consistent. A team builds an AI-generated insight they believe in, and then watches it stall, not because of a technical flaw, but because of organizational skepticism. The stakeholder on the receiving end of the insight does not trust it, cannot interrogate it, and is unwilling to stake their credibility on it. The conversations in the room sounded like this:

“ We vet every output by hand before it reaches the business.
Which kind of defeats the point.”

“ My job is on the line if this is wrong. So I double-check everything,
even when I know the model is right.”

This dynamic has a name worth using: the proof tax. It is the invisible cost of deploying AI in an organization that does not yet trust it. It shows up in the hours of manual validation, the line-by-line walkthroughs with skeptical stakeholders, and the cross-checks that must happen before any AI-generated insight is allowed to influence a decision. And the irony that surfaced repeatedly in the room is that the proof tax often does not just slow things down. It eliminates the productivity gain AI was supposed to deliver in the first place.

When it takes four hours to manually validate an insight the model produced in four minutes, the organization has not saved time. It has redistributed it.



Three Dynamics Worth Understanding

1. Trust Travels Slowly

The leaders in the room are not skeptics of AI. Most are genuinely invested in making it work and understand the potential clearly. The skepticism lives elsewhere, in the field sales rep who does not understand what a model is, in the marketing lead asking "how does it know that," and in the SVP who wants to know what happens if the organization acts on it and it turns out to be wrong.

Earning trust from a technically literate peer is manageable. Earning trust from a non-technical stakeholder who carries real accountability for the decisions that follow is a fundamentally different challenge. Most analytics and insights leaders are navigating that challenge without dedicated organizational support.

2. The Pilot Trap

A widely cited data point at the conference: approximately 5% of AI pilots scale to enterprise deployment. The room reflected this. Most leaders described their AI programs as internal, efficiency-focused, and still in limited deployment. The limiting factor was not the technology. It was the organizational infrastructure: governance frameworks, change management, stakeholder education, and the readiness of end users to trust what the models produce.

The aspiration across the industry is transformation. The reality, in most cases, is still augmentation. That gap is not primarily technical.

3. Asymmetric Accountability

Perhaps the most underappreciated dynamic in pharma AI right now is this: the political and career risk of a bad AI output is perceived as greater than the career reward of a good one. "You don't get fired for doing it the old way" remains the operating logic in many organizations. That logic is not irrational. It reflects how accountability for AI-assisted decisions is currently structured.

In private conversations during the conference, senior leaders from large pharma companies described the same pattern. One explained that even with an internal AI tool fully deployed, he personally reviews every output before it reaches anyone else because, in his words, "if garbage goes in and people interpret it in a busy environment, that's the scariest part." The loop is still manual. The trust is still personal.

Organizations that do not address this asymmetry will find their AI tools becoming permanent back-office utilities, useful internally but invisible in the decisions that matter, regardless of how capable the underlying technology is.

Broader Themes from Reuters Pharma USA 2026

These roundtable findings did not exist in isolation. They were consistent with themes running through the full two days at Reuters Pharma USA.

AI in life sciences is advancing and many of the solutions on display reflected real progress in making AI operational. At the same time, a quieter theme ran beneath the demonstrations. The industry is moving faster than its collective readiness. Terms like agentic systems, context windows, embeddings, and retrieval-augmented generation were used with varying degrees of shared understanding. The speed of deployment has outpaced organizational fluency, and that fluency gap introduces risk: teams deploying tools they cannot fully explain to the stakeholders they serve.

 *Speed of AI deployment is outpacing organizational fluency, and that fluency gap is introducing risk.*

A structural shift in decision-making was also evident. AI governance is centralizing. What used to be fragmented across business units is moving toward enterprise-level oversight, typically under Chief Data Officers (CDO) who are rebranding as Chief Data and AI Officers. That centralization changes the dynamics of adoption. The business champion who wants a tool is not the decision-maker. The technology and data leadership that controls the governance framework is. Enterprise deployment requires engaging both.





Glass Box Intelligence as a Trust Architecture

This is where the Glass Box Intelligence framework, introduced by Humigent in the roundtable, becomes more relevant. But for different reasons than originally anticipated.

Glass Box Intelligence is built on Humigent's core architectural conviction: when reasoning happens before the language model, through a Verifiable Logic Layer that validates data and applies explicit logic before any language is generated, the AI is not guessing. It is reporting on what verified reasoning found. That design addresses hallucination risk at its source, and it produces the traceable audit trail that regulators and compliance teams require.

When the framework was introduced in the roundtable, the pillar that generated the most engagement was the Verifiable Logic Layer. But the reason was not regulatory compliance. It was communication. Stakeholders who push back on AI outputs often do so not because they distrust the technology in principle, but because they have no way to follow the reasoning behind a specific output. A black-box result, even a correct one, asks them to extend a level of trust they are not prepared to give when their credibility is at stake.

Glass Box Intelligence changes that calculus. When the reasoning chain is explicit, inspectable, and traceable to specific data sources, the person delivering the insight has something they currently lack: a step-by-step answer to the question "how do you know?" That is not only an architecture for reliability. It is an architecture for organizational trust. And trust, as the room made clear, is the real bottleneck.



The value of auditable AI in pharma is not just satisfying a regulator. It is giving the person who has to defend the output something credible to stand behind.

This is a point worth dwelling on. Humigent's Glass Box Intelligence framework was designed around the conviction that AI in life sciences must be an operating model, not just a tool. The distinction matters. A tool produces output. An operating model defines how that output is generated, validated, explained, and ultimately trusted within the organization. When AI is treated as a tool, adoption depends on individual champions. When it is built as an operating model, adoption becomes structural.

A Three-Layer View of Pharma AI Maturity

One of the most useful frameworks to emerge during the conference came from a private conversation with senior leaders from two large pharma companies. They described three distinct layers of AI deployment in the industry today:

- **Layer 1: General-purpose AI** such as copilots, productivity tools, and chat interfaces. Most organizations have touched this layer. It is widely understood, relatively low risk, and increasingly commoditized.
- **Layer 2: Platform-embedded AI** built into tools organizations already use such as Veeva, Databricks, Snowflake, and Salesforce. This is where much of the current activity lives, often with limited integration across systems.
- **Layer 3: Enterprise-scale AI transformation** that fundamentally changes how organizations source insights, make decisions, and operate. This layer involves rearchitecting data infrastructure, harmonizing disparate data sources, and deploying AI across organizational boundaries with real governance, change management, and executive alignment.

Most of what was visible at Reuters Pharma USA 2026 fell into Layers 1 and 2. The aspiration, across nearly every organization represented, is Layer 3. The gap between those two things is where the real work lies, and it is not primarily a gap in technology. It is a gap in organizational readiness.



What This Means in Practice

The picture that emerged from the conference is one of an industry that sees AI's potential clearly but is still working through the organizational conditions required to act on it at scale. A few implications stand out.

AI literacy must extend beyond the analytics team. The fluency gap is real, not among the people building AI solutions, but among the stakeholders those teams serve. Field sales leaders, marketing VPs, and commercial executives are being asked to act on AI-generated outputs they do not fully understand. That gap will not close on its own. It requires deliberate education, framed not in technical terms, but in business outcomes.

Reasoning must be visible before the output reaches a skeptic. The most effective way to reduce the proof tax is not to ask stakeholders to trust AI more. It is to build systems where trust is easier to extend, where outputs arrive with clear provenance, explicit logic, and a confidence level the recipient can evaluate. That is a design choice, not a feature request.

Working product matters more than polished narratives. Governance is centralizing across pharma. When enterprise AI governance committees form, they tend to standardize on what is already working inside the organization. Organizations and their partners that have lightweight proofs of concept in place before those committees convene will have a structural advantage. The distinction is not whether to start small it is whether that starting point is designed with the operating model in mind.

AI must be an operating model, not an experiment. The organizations that will move from Layer 2 to Layer 3 are the ones building AI into their operational fabric, with measurable ROI, human-in-the-loop validation, and architecture that makes every output defensible. Treating AI as a series of disconnected pilots will produce interesting results. It will not produce transformation.



Looking Forward

The pharmaceutical industry is in the middle of a trust-building moment with AI. That is not a failure of the technology. It is a stage in its adoption. The skepticism heard in the roundtable and across the conference floor is not permanent. It reflects the natural friction that emerges when organizations encounter a powerful new capability before the infrastructure of understanding, accountability, and governance has caught up.

Humigent's position is that this readiness can be built, deliberately and systematically, through architecture that is reliable by design and auditable by nature, through human-in-the-loop workflows that keep the right people involved at the right moments, and through a commitment to making AI an operating model rather than a point solution.

The proof tax is real, but it is not inevitable. The organizations that invest in making their AI defensible, traceable, and understandable to the full range of stakeholders who must act on it will be the ones that move from experimentation to enterprise impact. The window for building that foundation is open now.

About Humigent

Humigent is a Human + AI Intelligence company empowering Life Sciences organizations to seamlessly adopt and scale Applied AI and Advanced Analytics. With more than 150 years of combined leadership experience across AI, Life Sciences, and Data Strategy, the Humigent team brings deep domain expertise, cutting-edge technology, and a flexible service model to some of the most complex challenges in the industry. Humigent's Glass Box Intelligence framework and R³ Cognitive Architecture (Retrieve, Reason, Respond) are designed for regulated environments where every output must be accurate, traceable, and defensible.

Learn more at humigent.ai | Contact: info@humigent.ai