From Chatbots to Colleagues:

Building Agentic AI for Biotech Research and Compliance

ptp Infinite Innovation | aws

A proteomics researcher at a mid-sized biotech company has a hypothesis about a transmembrane protein. Testing it in vitro would be a costly, months-long process. But testing it computationally could compress that timeline to days. Generative AI with reasoning capabilities could propose structural modifications and then validate them against a protein-folding model such as AlphaFold. The problem: the researcher's company has no approved AI tools, no governance policy, and no cloud infrastructure team that understands the life sciences stack.

That scenario captures the central problem facing biotechnology companies today. They are eager to participate in AI, particularly agentic AI. Salesforce research on AI in life sciences found that 94% of companies in this sector saw agentic AI as critical to scaling and strengthening operations.

Drug discovery is still a trial-and-error process in which researchers test successive proteins to find one that folds correctly and doesn't produce toxic effects. AI-powered agents that can chain tools together, reason through multi-step workflows, and incorporate feedback loops, could transform both the search and discovery problems. Despite its potential, deployment of this technology has been anything but straightforward. In its State of AI in Healthcare 2025 report, Menlo Ventures says that life sciences companies are relatively early on in the AI adoption curve.

The challenges are well-documented but poorly addressed. Life sciences companies face fragmented data, missing governance frameworks, and conservative IT policies. There's also a skills gap that leaves those who most need AI unable to provision its underlying infrastructure.

In this playbook, we'll examine those challenges and explore a security architecture that enables adoption without sacrificing compliance. We'll investigate the acceleration strategies that work in practice and the real-world deployments that demonstrate what's achievable when the pieces come together.

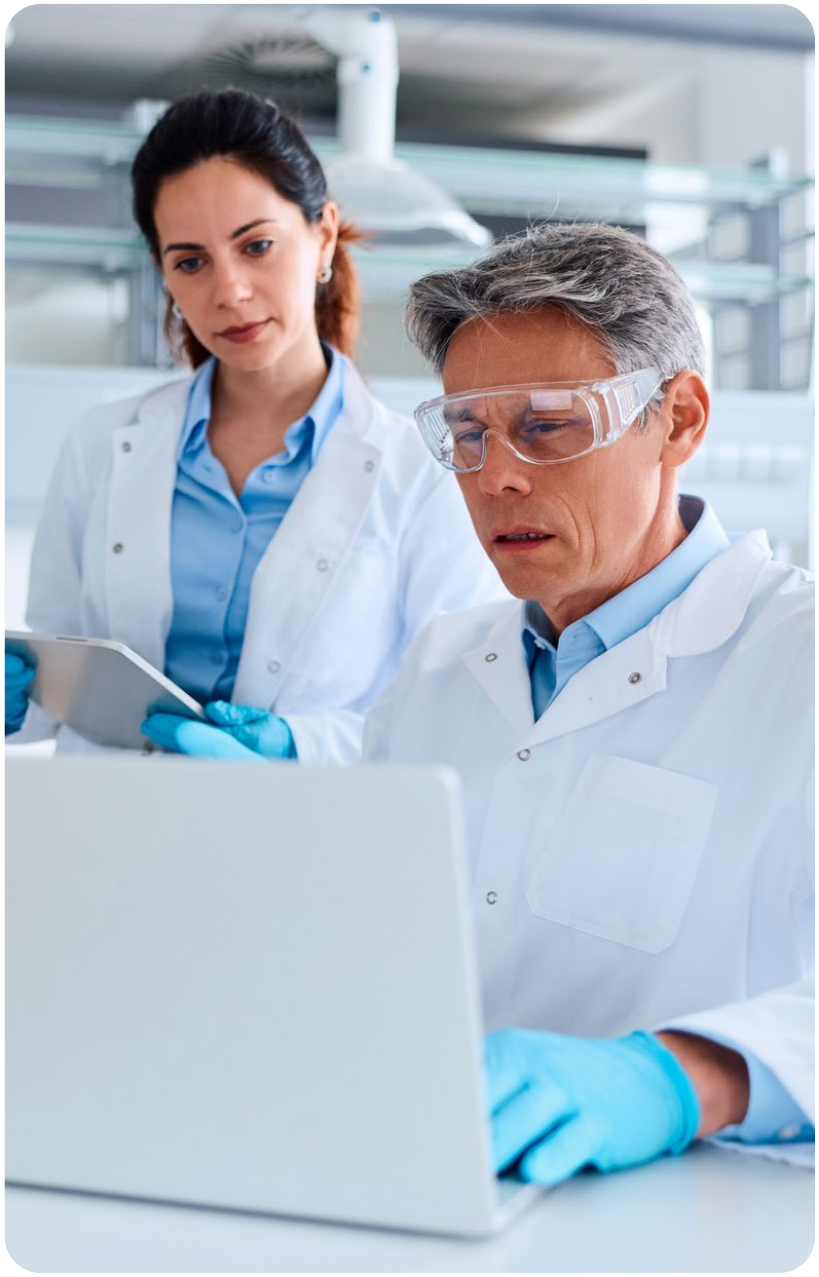

ptp Infinite Innovation

The Challenges of Agentic AI in Life Sciences

When we break down the hurdles facing life sciences companies in AI, the data problem comes first. Everything else depends on it. Biotechnology data lives in a series of silos, ranging from laboratory information management systems to electronic lab notebooks, clinical trial databases, regulatory submission archives, and simple email threads. Often, there is no integration layer connecting them.

DROWNING IN DATA

Fragmentation is just one challenge. The sheer volume dwarfs what most organizations have the tooling to process. Many life sciences companies haven't completed basic business intelligence work on their existing datasets, so an agentic AI pilot would be equivalent to running before they can walk.

A LACK OF GOVERNANCE GUIDELINES

That immaturity compounds a governance problem that may be even more damaging. Organizations frequently lack formal guidance on which AI tools their staff may use, what data those tools may access, and how to protect intellectual property when interacting with external models.

In a sector governed by HIPAA, GDPR, and FDA submission requirements, the absence of a clear policy makes even the most cautious experimentation difficult. Researchers might sign up for consumer AI services with corporate credit cards and feed them proprietary data, unaware that the model providers may use it for retraining. They incur a compliance exposure that nobody tracks until IT discovers the charges and shuts everything down.

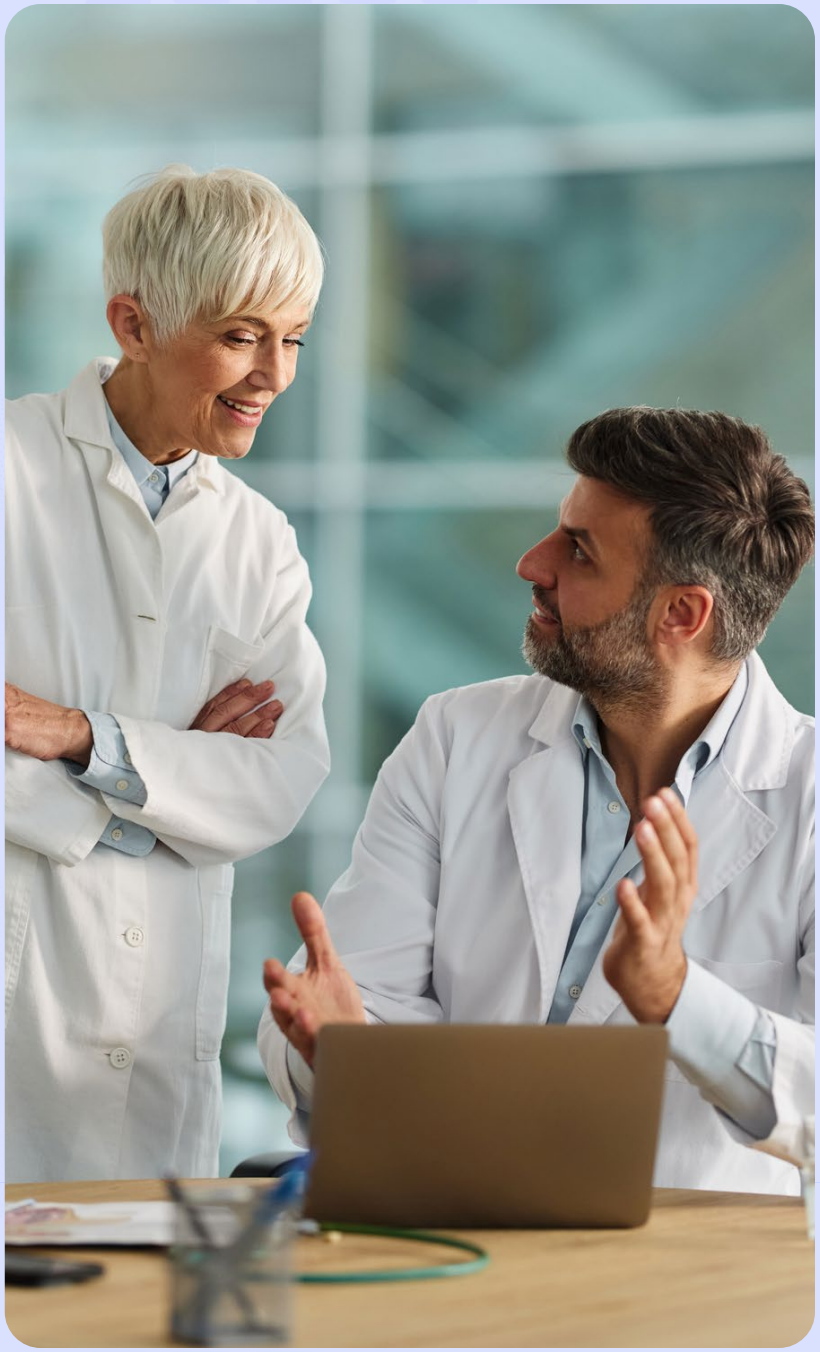

MISMATCHED SKILLS

Even if there were clear guidelines from IT on what to use, life sciences companies often lack the provisioning frameworks to enable researchers to build the necessary infrastructure. Data scientists who can write fluent Python are unlikely to know how to provision an EC2 instance on AWS. They understand their experimental domains with extraordinary precision, but they are not DevOps experts. Telling someone to launch a virtual GPU instance with all the right AI tooling is like telling a concert musician to build a piano.

The practical implication is stark: either organizations invest in training researchers to manage cloud infrastructure, or they abstract the technology away entirely, giving scientists tools that hide the plumbing and let them focus on asking the right questions. The second path is faster, more realistic, and less likely to produce the kind of bespoke, unsupported infrastructure that falls apart when its creator leaves the company.

MISMATCHED MODELS

One more barrier deserves attention, because it affects the most ambitious use cases. While agentic AI models work well for broad information search and document preparation, the specialized machine learning models needed to validate generative work, such as protein structure prediction and molecular interaction modeling, remain immature in many subdomains. A team working on transmembrane proteins, for instance, may find that the best available validation model is only 30–40% accurate, roughly where AlphaFold was in 2017. The tools for generating hypotheses have outpaced the tools for testing them, and that gap limits what agentic workflows can reliably accomplish today. The Menlo Ventures report says that life sciences companies are still mostly using general-purpose LLMs. However, two-thirds would like to build or fine-tune proprietary models, it adds.

Creating Security That Doesn't Strangle Innovation

The instinct to lock down AI access is understandable in a regulated industry, but it misdiagnoses the problem. Researchers don't use unauthorized tools because they're reckless; they use them because approved alternatives don't exist. Building those alternatives requires a governance framework that starts at the executive level and flows down through every function that touches AI. That ranges from compliance through IT and security to laboratory science and legal.

The framework must address three questions:

- **What data classification scheme applies to AI interactions?**
  Not all data carries the same risk. A model trained on publicly available research papers raises different concerns than one trained on patient-identifiable clinical trial records.
- **How does intellectual property protection work when proprietary sequences or molecular structures enter an AI system?**
  The answer depends on whether data remains within the organization's infrastructure boundary or leaves it.
- **Who audits AI-generated outputs before they enter regulatory submissions or influence experimental design?**
  AI governance without human review checkpoints isn't governance so much as automation of liability.

The architectural choices that satisfy these requirements tend to converge on a common pattern: enterprise AI platforms that keep data within the organization's cloud environment, support fine-grained access controls, and maintain audit logs sufficient for regulatory review.

A platform for clinical AI should provide model access without transmitting training data outside the customer's account. That's a meaningful security distinction compared to consumer AI services, where data-handling policies are opaque or subject to change. The platform matters less than the principle: data sovereignty, auditability, and the separation of model inference from model training.

Getting security right unlocks everything else. Without it, every AI project risks the same fate: a promising pilot, followed by a compliance review and a shutdown.

How to Accelerate Development

The most important thing biotech teams misunderstand about agentic AI is its scope. Conversational interfaces (chatbots that answer questions about internal documents) represent the simplest deployment pattern, and many organizations stop there. But effective AI implementation extends well beyond conversation. It uses agent chains, tool orchestration, and iterative feedback mechanisms that allow systems to plan and execute multi-step workflows, evaluate the results, and refine their approach.

A chatbot can search a laboratory information system and return results. An agent chain can search that system, cross-reference findings against published literature, identify experimental protocols that match a researcher's current hypothesis, flag potential regulatory conflicts, and draft a summary document in the organization's preferred format - all from a single natural-language request. Jumping from the first capability to the second changes the relationship between researcher and tool. AI then functions as a colleague rather than a mere search engine.

Finding the right use cases matters more than finding many of them. Organizations that form committees to brainstorm AI applications tend to generate long lists of marginal improvements. Those that look at real patterns in researcher demand will be better equipped to prioritize cases. To surface those patterns, explore which questions researchers ask most frequently, which workflows consume disproportionate time, and which bottlenecks delay regulatory submissions.

Here are some likely use cases for agentic AI in a life sciences environment:

CANCER BIOMARKER DISCOVERY THROUGH MULTI-AGENT ORCHESTRATION

A researcher hunting for cancer biomarkers typically pulls clinical and RNA-seq data from a database using SQL, runs survival analyses in a stats package to generate Kaplan-Meier plots, and searches PubMed and internal repositories to see whether the finding is already known. That's a time-consuming process. A supervisor agent does all of that on its own, handing each step to a specialist sub-agent and pulling the results back into a single answer with citations, freeing up the researcher's time for other work.

GENOMIC VARIANT INTERPRETATION AT PIPELINE SCALE

Interpreting genetic variants typically involves running an annotation job for each patient file using a tool like ClinVar. The research must then wait for the jobs to finish, check the results against clinical significance databases, and write up what they found. This is something an agent can handle from end to end. Conversely, a plain chatbot can't because it only responds in the moment and can't manage work that runs in the background.

MULTIMODAL PATIENT DATA CORRELATION FOR TRANSLATIONAL RESEARCH

A translational researcher normally has to open multiple images, such as CT scans, whole-slide pathology images, and X-rays, in different tools, mentally stitching the findings together with information from other clinical records. Conversely, a coordinating agent can send each file type to a specialist sub-agent and combine the results.

CLINICAL TRIAL PROTOCOL GENERATION GROUNDED IN PRECEDENT

Someone drafting a new trial protocol normally searches ClinicalTrials.gov for similar studies, reads through their eligibility criteria and endpoints to see what worked, then writes the new protocol from scratch against the Common Data Model. An LLM agent can retrieve the precedents and produce a first draft with inclusion criteria, endpoints, and a statistical plan already filled in, streamlining the process.

BIOMEDICAL RESEARCH AGENT WITH 30+ BIOMNI DATABASE TOOLS

A researcher seeking answers to a question about a protein or disease pathway typically has to know which database holds the answer and query each database separately. Agentic AI can leverage an LLM's reasoning to select the appropriate database for the question and then pull the answers together.

AGENTIC ANALYSIS OF FDA SUBMISSIONS

In this real-world case, a biotech company needed a consistent voice and formatting across FDA submissions prepared by dozens of researchers. It trained an LLM on 380,000 documents from the CEO's communications archive, then built templates for FDA documentation in the CEO's preferred format. Researchers can now generate late-stage drafts that maintain organizational consistency across all submissions, reducing the number of revision cycles needed before peer review.

Working with a partner that has agents such as these preconfigured helps to collapse the timeline from concept to deployment. It also demonstrates value to leadership teams that are still skeptical about AI's practical applications.

What Effective AI Implementation Partnerships Look Like

Most biotech organizations lack the staff to architect, deploy, and maintain enterprise AI infrastructure while simultaneously managing the regulatory requirements that govern their industry. That gap defines what a successful implementation partnership should deliver: not product sales, but the combination of cloud infrastructure expertise and life sciences domain knowledge that most organizations can't assemble internally.

A typical engagement involves three parallel workstreams:

- **Building AI governance frameworks** that satisfy both internal security teams and external auditors.
- **Connecting AI systems** to legacy laboratory and clinical databases where an organization's most valuable data already lives.
- **Training bench scientists** on the approved tools.

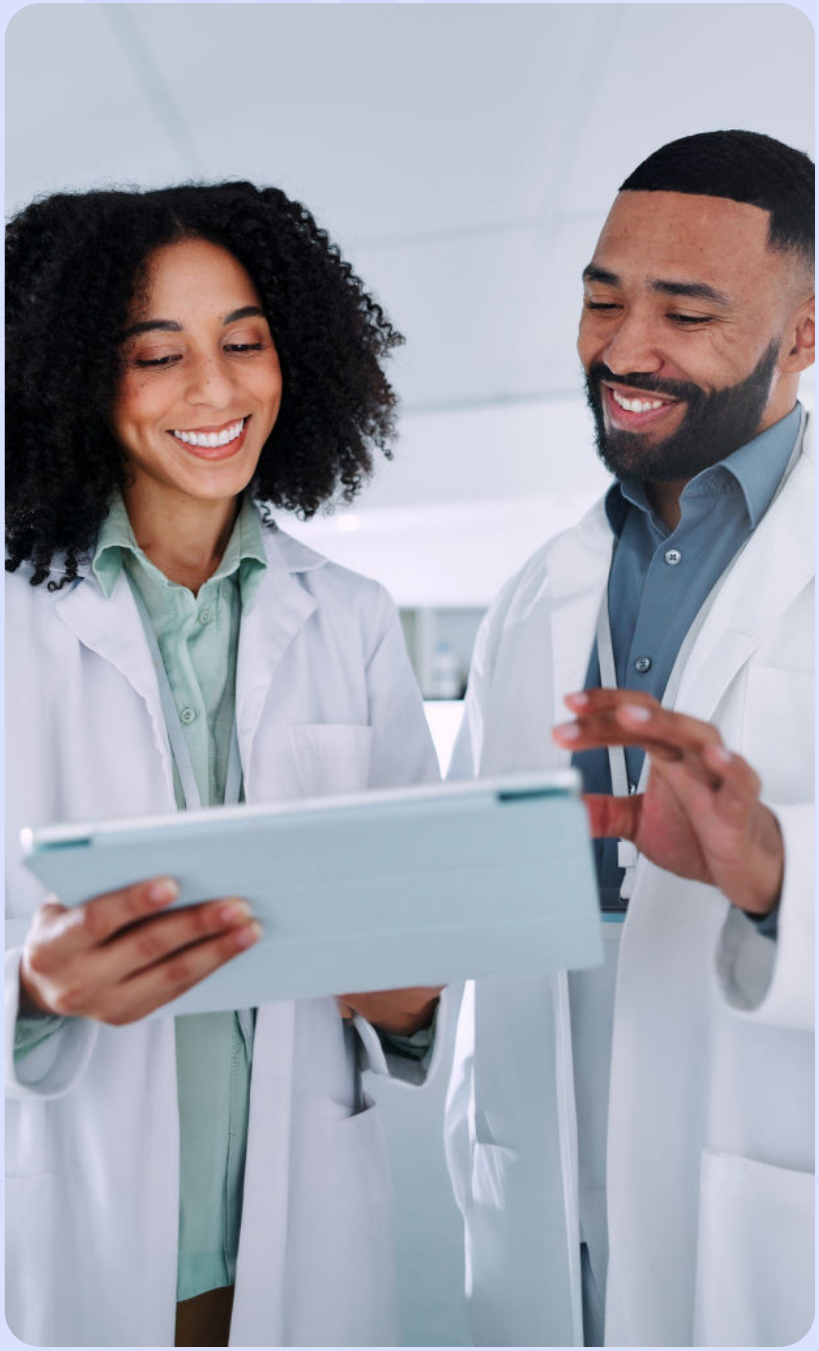

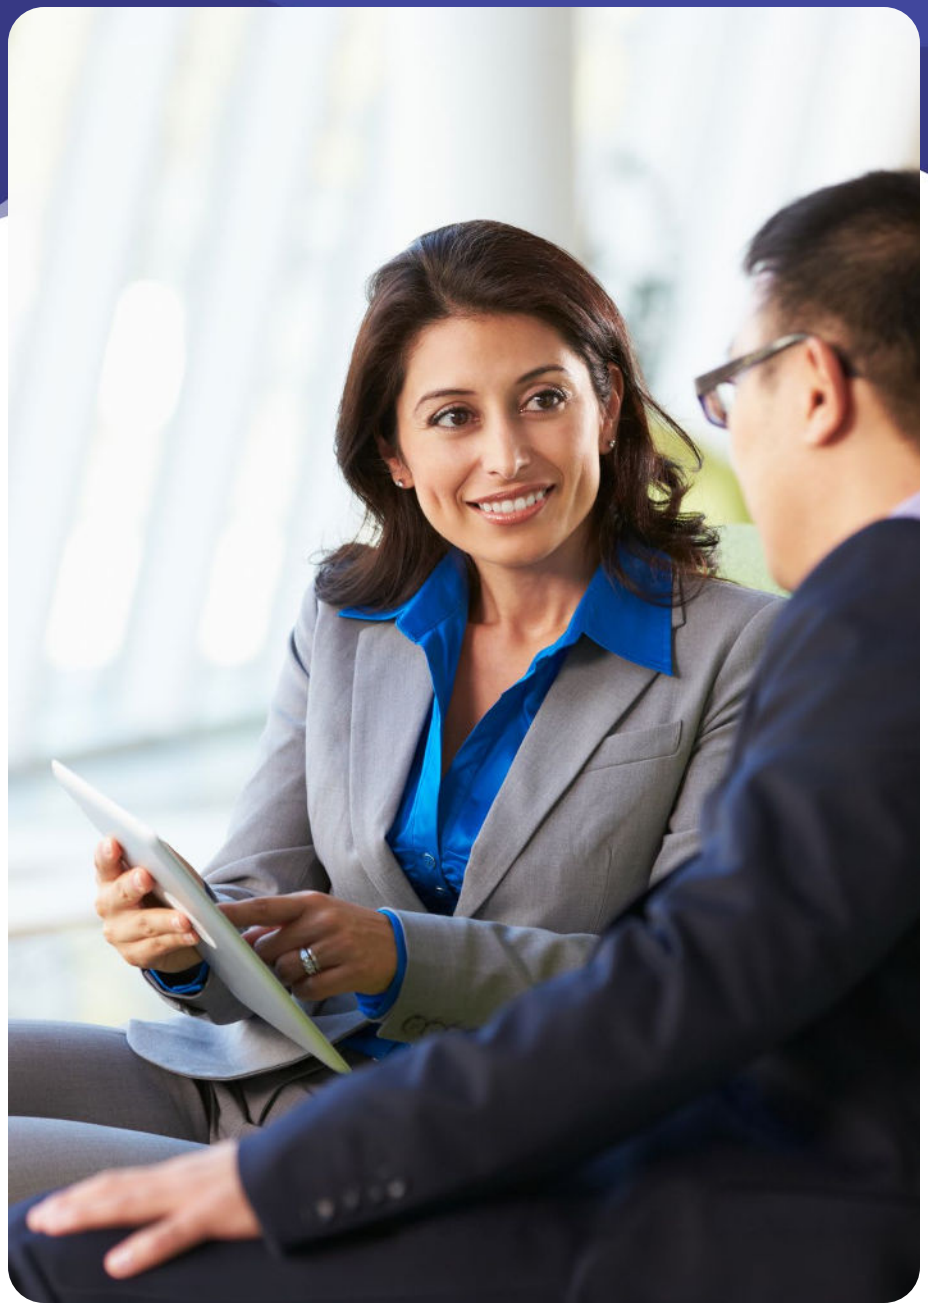

Conclusion

Biotech organizations that establish governance frameworks and deploy operational AI tools now will compound their advantages over those still deliberating. Every month spent in policy limbo is a month where researchers use unauthorized consumer tools or don't benefit from AI at all.

Rapid adoption requires two things that only executive leadership can provide: a formal governance structure and a cultural willingness to experiment within defined boundaries. The most effective approach is a central, multi-disciplinary committee with explicit responsibility for the organization's entire AI initiative. That can develop a coordinated strategy spanning compliance, IT, security, and laboratory science.

A committee-led approach defines what data AI may access, which tools are approved, how outputs are reviewed, and where the organization invests next. Without that authority, governance remains reactive and innovation stays constrained. A concerted approach to AI in life sciences will narrow the gap between what's possible and what's permitted.

PTP exists to accelerate scientific innovation by solving the complex through technology & services that are intelligent, secure and frictionless. Purpose-built for biotechnology environments, we combine deep scientific understanding with advanced technology management across CloudOps, SecOps, NetOps, and UserOps. Our experts support life sciences organizations through their complete lifecycle—from preclinical discovery to global commercialization. By managing technical complexity and ensuring continuous compliance, we enable research teams to focus entirely on advancing treatments from discovery through commercialization.

Learn more at ptp.cloud.com.

Amazon Web Services (AWS) is guided by customer obsession, pace of innovation, commitment to operational excellence, and long-term thinking. By democratizing technology for nearly two decades and making cloud computing and generative AI accessible to organizations of every size and industry, AWS has built one of the fastest-growing enterprise technology businesses in history. Millions of customers trust AWS to accelerate innovation, transform their businesses, and shape the future. With the most comprehensive AI capabilities and global infrastructure footprint, AWS empowers builders to turn big ideas into reality.

Learn more at aws.amazon.com and follow @AWSNewsroom.

Studio
by informa techtarget

Expert led.
Impact driven.

Studio is Informa TechTarget's global content studio offering brands an ROI rich tool kit: Deep industry expertise, first-party audience insights, an editorial approach to brand storytelling, and targeted distribution capabilities. Our trusted in-house content marketers help brands power insights-fueled content programs that nurture prospects and customers from discovery through to purchase, connecting brand to demand.

Learn more