IntuitionLabs

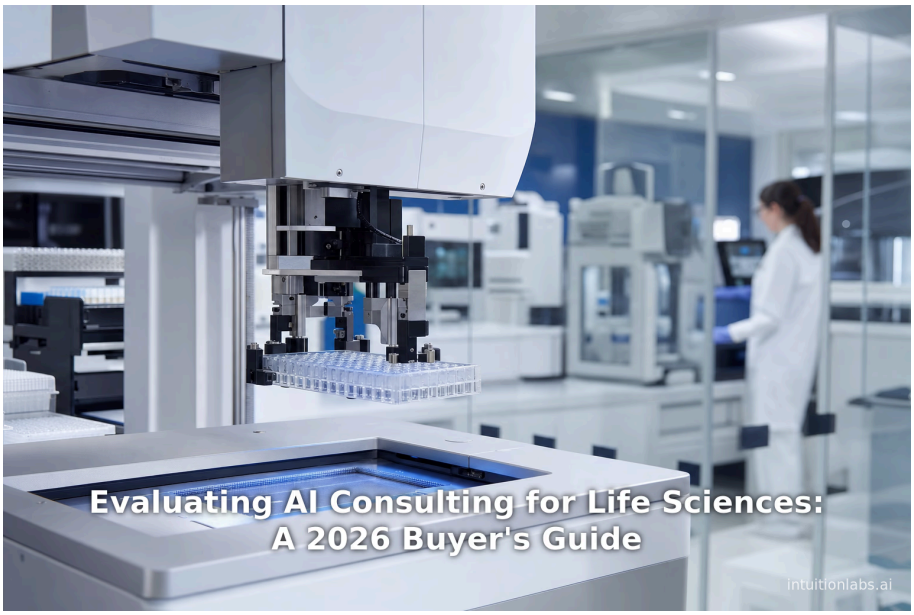
INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Evaluating AI Consulting for Life Sciences: A 2026 Buyer's Guide

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-06-30 · ai consulting for life sciences · life sciences ai consulting · pharmaceutical ai consulting · biotech ai consulting · ai consulting firms for pharma · how to choose an ai partner · evaluating ai vendors life sciences · artificial intelligence life sciences consulting

Original article: https://intuitionlabs.ai/articles/evaluating-ai-consulting-for-life-sciences



In this article

1. Executive Summary
2. Introduction and Background
3. The AI Consulting Landscape: Global Integrators, Boutique Specialists, and Platform Partners
4. Regulatory Foundations That Shape Every Engagement
5. Data Governance, Privacy, and Vendor Evaluation Frameworks
6. Implementation Considerations and Process Changes
7. Data Analysis and Evidence
8. Case Studies and Real-World Examples
9. Pfizer and Amazon Web Services: The PACT Model for Rapid, Governed Prototyping
10. Novartis and Microsoft: A Five-Year Enterprise AI Alliance
11. Insilico Medicine: The First AI-Designed Drug to Approach Phase 3
12. IntuitionLabs and Scilex Holding Company: Commercial Analytics via Certified Veeva Implementation
13. Implications and Future Directions
14. Frequently Asked Questions (FAQs)
15. Conclusion

Executive Summary

Life sciences companies are moving artificial intelligence (AI) from pilot to production faster than most other regulated industries, and the market for outside help is scaling with it. The **AI in life sciences market** was valued at **$3.61 billion in 2025** and is projected to reach **$13.64 billion by 2031**, a **24.78% compound annual growth rate (CAGR)** (mordorintelligence.com). The broader **AI consulting services market** across all industries was valued at **$22.27 billion in 2025**, heading toward **$349.80 billion by 2034** at a **35.8% CAGR** (marketdataforecast.com). Deloitte's web summary states that the life sciences industry is poised to potentially unlock **$5 billion to $7 billion in value** through artificial intelligence (AI), particularly GenAI; the underlying report bases that figure on peak value for a single top-10 biopharma company with average revenue of $65 to $75 billion, captured over five years (deloitte.com), with research and development (R&D) accounting for **30% to 40%** of the potential value and commercial applications another **25% to 35%** (deloitte.com).

Demand is being pulled by real, quantified results as much as by hype. NVIDIA's 2025 survey of more than 600 healthcare and life sciences professionals found **81% report AI has helped increase revenue**, and **69% of pharmaceutical and biotech respondents are actively using Generative AI**, second only to digital healthcare among the segments surveyed (blogs.nvidia.com). Pfizer's collaboration with Amazon Web Services (AWS), the Pfizer-Amazon Collaboration Team (PACT), is estimated by Pfizer to be able to save scientists up to **16,000 hours of search time annually** and cut infrastructure costs by **55%** (aws.amazon.com). Insilico Medicine's AI-discovered idiopathic pulmonary fibrosis (IPF) candidate, rentosertib, entered a Phase 3 trial in 2026, with its first patient dosed in the GENESIS-IPF-3 study in September 2026 (insilico.com), though it still faces the same **55% likelihood of Phase 3 success and multi-hundred-million-dollar cost** that governs any traditional drug candidate (chemistryworld.com).

But adoption success is far from guaranteed, and this is precisely where evaluating an AI consulting partner well or poorly determines outcomes. A widely cited 2025 Massachusetts Institute of Technology (MIT) NANDA report, based on 150 leader interviews and analysis of 300 public AI deployments, found that **95% of enterprise GenAI pilots fail to deliver measurable profit-and-loss impact** (fortune.com).

Regulatory scaffolding is also tightening the field for consultants who lack genuine domain depth. The **U.S. Food and Drug Administration (FDA)** published its first draft guidance on AI in drug and biological product development in January 2025, introducing a risk-based credibility assessment framework tied to a specific "context of use" (fda.gov). The **European Union (EU) AI Act**'s high-risk obligations phase in on a staggered timeline running through **2 December 2027 and 2 August 2028** (digital-strategy.ec.europa.eu), and the International Society for Pharmaceutical Engineering (ISPE) has published a dedicated Good Automated Manufacturing Practice (GAMP) 5 guide on validating AI-enabled GxP systems (ispe.org). This report, prepared with sources current through **September 2026**, examines how to evaluate AI consulting for life sciences: the market landscape, the major consulting archetypes (global systems integrators, boutique domain specialists, and technology-platform partners), the regulatory and data-governance constraints unique to the industry, the economics of engagement pricing, and the evidence from named deployments, both alliances and implementation partners. Life sciences and AI advisory firms such as **IntuitionLabs**, the publisher of this article and a life-sciences AI consultancy that describes itself as a Veeva X-Pages Partner, argue (per intuitionlabs.ai's own site) that regulatory-native delivery, not generic AI fluency, is what should differentiate a shortlist, a view this report evaluates against the broader competitive and regulatory evidence throughout.

Introduction and Background

Artificial intelligence consulting for life sciences sits at the intersection of two demanding disciplines: enterprise AI implementation, which is hard everywhere, and pharmaceutical and biotechnology regulation, which tolerates almost no ambiguity. A pharmaceutical or biotechnology organization considering an AI initiative in drug discovery, clinical operations, regulatory affairs, manufacturing, or commercial operations faces a fundamentally different risk calculus than a retailer deploying a recommendation engine: a poorly validated model touching a Good Clinical Practice (GCP), Good Manufacturing Practice (GMP), or Good Laboratory Practice (GLP) process, collectively known as **GxP**, can trigger findings from the FDA, the European Medicines Agency (EMA), or another regulator, and a hallucinated citation in a regulatory submission or medical information response carries patient-safety consequences that a marketing chatbot error does not.

Survey data confirms that pharmaceutical leadership treats this as urgent. ZS Associates' 2025 survey of 127 pharmaceutical, biotechnology, and life sciences technology executives found **93% anticipate increased investment** in data, digital, and AI for the year, up from 88% the year prior (zs.com). Yet the same survey found that only **around 10% of employees** with company-approved access to Generative AI tools use them on a weekly basis, and **71% of respondents** think it is time to moderate GenAI investment pace given the nontechnical work still required to realize value (zs.com). That gap between budget intent and adoption maturity is exactly the gap AI consultants are hired to close, and exactly where vendor selection quality determines whether the engagement produces a validated, auditable system or an expensive prototype that never reaches production.

"AI consulting for life sciences" is not a single service category. It spans four overlapping archetypes: **global systems integrators and strategy houses** (Accenture, Deloitte, McKinsey, L.E.K.) that bring scale, change-management muscle, and hyperscaler alliances; **boutique life sciences AI specialists** (Statsby Solutions, Sciagen AI, Phintel, Siglasciences) that trade scale for narrower, deeper domain focus; **platform and implementation partners** tied to a specific enterprise system, most notably the Veeva ecosystem that underpins most large pharmaceutical commercial and regulatory operations; and **adjacent advisory firms**, exemplified by IntuitionLabs, that combine AI and analytics advisory with hands-on platform implementation rather than pure strategy slides. This report evaluates all four archetypes against the criteria that matter most in a regulated environment: regulatory fluency, data governance rigor, proof of production (not just pilot) deployments, and pricing transparency. It also situates the discussion within the current wave of FDA, EMA, and EU regulatory activity that is reshaping what "compliant AI" even means as of late 2026.

The AI Consulting Landscape: Global Integrators, Boutique Specialists, and Platform Partners

Global systems integrators dominate AI consulting revenue but not necessarily life sciences AI depth. **Accenture** reported **$5.9 billion in Generative AI new bookings for fiscal year 2025**, alongside **$80.6 billion in total new bookings** and 7% overall revenue growth (newsroom.accenture.com). Accenture's life sciences practice cites survey data claiming **100% of biopharma executives** agree AI foundation models will "revolutionize" how AI is used across data types, and **92%** say technology is critical to reinvention strategy (accenture.com). The same survey found **90% of biopharma executives** agree that data transparency is becoming a competitive differentiator, and **60% of life sciences supply chain executives** want better real-time inventory and visibility from external manufacturing partners (accenture.com). Other strategy houses with dedicated life sciences practices include **L.E.K. Consulting**, whose research projects that roughly **50% of future biopharma revenue growth** will come from external innovation such as mergers, acquisitions, and licensing rather than internal R&D (lek.com), and **Clarkston Consulting**, which spans data and analytics, digital and technology, quality and compliance, and supply chain advisory specifically for life sciences clients (clarkstonconsulting.com). Deloitte, similarly, frames its life sciences AI practice around the **$5 billion to $7 billion value opportunity** referenced above (a peak-value estimate for a single large biopharma company, not an industry total), spanning 20 AI use cases across R&D, manufacturing, and commercial operations, with an estimated **five-year timeline** for an enterprise to capture peak value (deloitte.com). These firms offer breadth: global delivery capacity, established hyperscaler alliances (Accenture partners with Google, Amazon, Microsoft, and Veeva Systems, among others, to "pharmatize" AI, according to EVERSANA's description of the broader partner ecosystem) (eversana.com), and enough balance-sheet weight to absorb liability on very large transformation programs. Independent analyst recognition reinforces the scale argument: Accenture states it has been "recognized as Leader in Life Sciences R&D AI in Clinical Trials 2025" and separately positioned as the top-designated Leader on both the Vision and Capability and Market Impact axes of Everest Group's Veeva Services PEAK Matrix Assessment 2025 (accenture.com). Accenture's life sciences consulting role requirements include experience in research and development, clinical operations, quality, regulatory, or

related areas, along with practical knowledge of GLP, GCP, GMP, and validation processes (accenture.com). Separately, published MBB-tier billing estimates are roughly **$400 per hour for associates and $1,000-plus per hour for partners, inclusive of firm overhead** (core-mba.pro).

Boutique life sciences AI specialists position themselves explicitly against that gap. **Statsby Solutions** markets itself as "domain-native, production-grade, regulatory-aware," spanning clinical and enterprise data platforms, Generative AI, agentic automation, and Machine Learning Operations (MLOps) specifically for pharma, biotech, and clinical research organizations (statsby.ai). **Sciagen AI** describes itself as "practitioners led, not consultant led," staffed by people who have worked inside clinical operations, regulatory affairs, quality, safety, and commercial roles rather than career strategists (sciagen.ai). **Phintel**, an AI-first firm founded by former McKinsey and BCG consultants, combines human expertise with autonomous agents but stays deliberately small, with a LinkedIn-listed headcount of **1 to 10 employees**, executing each engagement with just one or two seasoned industry experts augmented by AI tooling (linkedin.com/company/phintel). Phintel describes its own model as combining "human expertise with AI and autonomous agents to help biotech and pharma leaders make sharper R&D, portfolio, and partnering decisions," a positioning that only works if the firm's small headcount is offset by genuinely deep individual domain expertise rather than raw delivery capacity (linkedin.com/company/phintel). This trades scale for depth: a boutique cannot staff a 200-person global rollout, but it typically cannot be fooled by a vendor demo that looks impressive to a generalist evaluator either. **iuvo Technologies** makes a similar case for IT-rooted pharma consulting, framing its pitch around the observation that "implementing [AI] within a regulated framework requires precision, validation, and deep industry knowledge" (iuvotech.com), while **VioSenza** occupies a related but distinct niche, embedding senior industry leadership directly inside client organizations on 3, 6, or 12-month retainers across seven practice areas including Technology and AI, rather than delivering a traditional project-based engagement (viosenza.com). **Siglasciences** distills the governance argument to a single line worth any buyer repeating back to a prospective vendor during due diligence: "If your AI can't explain itself to an auditor, it can't be deployed" (siglasciences.ai).

Platform and implementation partners occupy a third category, tied to the specific enterprise systems that already run most regulated pharmaceutical commercial and quality operations, chiefly the **Veeva** suite (Vault CRM, Vault eTMF/RIM/QMS/Safety, Nitro, and the newer **Vault CRM X-Pages** framework that replaced the earlier MyInsights product). Veeva formally certifies partners through its **X-Pages Certification Program**, requiring completed specialized training and demonstrated best-practice competency before a firm can be listed (veeva.com). Because Vault CRM X-Pages is described by Veeva as a "no-code, AI-enabled platform for custom field analytics," implementation partners in this category increasingly deliver AI capability directly inside the system of record rather than as a bolt-on layer, which materially reduces integration risk relative to a stand-alone AI tool that has to be wired into Veeva after the fact.

Table 1 below summarizes the primary evaluation dimensions a life sciences buyer should apply across these consulting archetypes.

Evaluation Dimension	What to Verify	Why It Matters in Life Sciences
Regulatory fluency	Named engagements referencing FDA, EMA, or GxP frameworks; familiarity with the FDA's January 2025 AI credibility guidance and the ISPE GAMP AI guide	Generic AI consultants rarely have hands-on validation-package experience; a mis-scoped "context of use" can invalidate an entire submission strategy (fda.gov)
Data governance depth	Explicit HIPAA Safe Harbor/Expert Determination and GDPR anonymization methodology, not just a mention of "compliance"	HHS defines two legally distinct de-identification standards; a consultant who conflates them creates re-identification risk (hhs.gov)
Production track record	Case studies with named clients, measured outcomes, and post-go-live support, versus proof-of-concept-only portfolios	95% of enterprise GenAI pilots fail to scale to measurable P&L impact, per MIT NANDA research (fortune.com)
Pricing transparency	Whether fee structure is disclosed (hourly, retainer, fixed-fee, or outcome-based) and benchmarked against market rates	Indicative web estimates (not specific to life sciences AI) put independent strategy consulting at $150 to $500 per hour, while MBB-tier billing runs roughly $400 associate/$1,000-plus partner per hour; misaligned engagement models create scope disputes (consultfees.com) (core-mba.pro)
Platform-native delivery	Whether the vendor is a certified implementation partner of the systems already in production (e.g., Veeva's X-Pages Certification Program)	Certified partners have completed formal training and demonstrated competency, reducing integration and upgrade risk (veeva.com)

This comparison does not rank vendors against one another; it lays out the criteria a life sciences buyer should demand evidence for, regardless of which archetype a shortlisted firm falls into. Global integrators tend to score highest on scale and lowest on unit-cost transparency; boutiques often invert that profile; platform partners score highest on integration risk reduction but narrowest on strategic scope.

Regulatory Foundations That Shape Every Engagement

No AI consulting engagement in life sciences can be evaluated apart from the regulatory environment it must operate inside, and that environment has moved substantially since the start of 2025. In January 2025, the FDA published its first draft guidance specifically addressing AI use in the development of drugs and biological products, titled "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products." The guidance establishes a **risk-based credibility assessment framework** for establishing and evaluating the credibility of an AI model for a particular "context of use" (COU), meaning the specific role and scope in which the model's output will be relied upon (fda.gov). A consultant advising on any AI system that touches drug safety, efficacy, or quality data needs fluency in this COU framework, because scoping it incorrectly at the start of a project can force a costly re-validation later. Internally, the FDA's Center for Drug Evaluation and Research (CDER) established a dedicated **CDER AI Council in 2024** to provide

oversight, coordination, and consolidation of the agency's AI activities, and CDER officials collaborated with the EMA to develop **10 guiding principles of good AI practice** that industry can reference when using AI to accelerate drug and biological product development (fda.gov).

The EMA has pursued a parallel but distinct track. Its **reflection paper on the use of AI in the lifecycle of medicines**, developed jointly with the Heads of Medicines Agencies-EMA (HMA-EMA) Big Data Steering Group, addresses AI applications from drug discovery through post-authorization pharmacovigilance, and explicitly calls for a "human-centric approach" to guide all AI development and deployment (ema.europa.eu). Where an AI or Machine Learning (ML) system is expected to affect the benefit-risk balance of a medicine, the EMA advises developers to seek early regulatory support through qualification of innovative development methods or formal scientific advice, a step many first-time AI vendors are unaware even exists (ema.europa.eu). The reflection paper itself sits inside a longer-running program: it forms part of the **Workplan 2022-2025 HMA-EMA joint Big Data Steering Group** initiative to build the European medicines regulatory network's capability in data-driven regulation, developed jointly with EMA's Committee for Medicinal Products for Human Use (CHMP) and its Committee for Veterinary Medicinal Products (CVMP) (ema.europa.eu).

The broadest and most consequential regulatory shift, however, is the **EU AI Act**. It entered into force on **1 August 2024** and applies in stages, with several deadlines postponed by the AI Omnibus, which was formally adopted as Regulation (EU) 2026/1744 and is now in force (eur-lex.europa.eu). Prohibited AI practices and AI-literacy obligations already applied from 2 February 2025, and governance rules for general-purpose AI (GPAI) models became applicable 2 August 2025 (digital-strategy.ec.europa.eu). Under the AI Omnibus (Regulation (EU) 2026/1744), which followed a political agreement reached on **7 May 2026**, the rules governing high-risk AI systems used in areas such as biometrics, critical infrastructure, employment, and border control apply from **2 December 2027**, while rules for AI systems embedded into regulated products (a category that will typically capture AI-enabled Software as a Medical Device, SaMD) apply from **2 August 2028** (digital-strategy.ec.europa.eu). The extended transition period is now written into the amended AI Act, and any consultant quoting the original high-risk deadlines without acknowledging the omnibus revision is working from stale information.

Complementing region-specific regulators, ISPE's **GAMP 5 Guide: Artificial Intelligence** provides the industry-standard risk-based validation methodology for AI-enabled computerized systems operating in GxP areas. ISPE describes the guide as "the single source for a holistic interpretation on effectively developing and using AI-enabled computerized systems in GxP areas, while safeguarding patient safety, product quality, and data integrity," and it defines activities and responsibilities for both internal teams and external suppliers across the AI system life cycle (ispe.org). Consultants unable to walk a client team through GAMP 5's risk-based validation categories in a discovery call are, at minimum, not equipped for GxP-adjacent AI work, regardless of how strong their general AI engineering may be.

Data Governance, Privacy, and Vendor Evaluation Frameworks

Because life sciences AI systems routinely touch protected health information (PHI) and personal data, any consulting evaluation must scrutinize data governance methodology specifically, not treat "HIPAA and GDPR compliant" as a checkbox. The U.S. Department of Health and Human Services (HHS) recognizes exactly **two legally valid de-identification methods** under the HIPAA Privacy Rule: **Expert Determination**, in which a qualified statistical expert certifies that re-identification risk is very small, and **Safe Harbor**, which requires removal of 18 specific identifier categories (hhs.gov). A consultant proposing to train or fine-tune a model on clinical or patient data should be able to state, in specific terms, which of these two paths a given dataset uses and why, since the methods carry materially different risk profiles and are not interchangeable. The Safe Harbor method's identifier list is granular enough to trip up an AI pipeline built without regulatory input: geographic subdivisions smaller than a state must generally be removed, and even ZIP codes require the first three digits to be changed to "000" whenever the corresponding geographic unit contains **20,000 people or fewer**, a threshold that a data engineer unfamiliar with HIPAA could easily miss when feature-engineering a training set (hhs.gov).

Beyond privacy mechanics, procurement teams increasingly formalize AI vendor selection through weighted scorecards rather than informal impressions. A structured framework typically scores candidate vendors across seven criteria categories: **technical capability, regulatory compliance, data handling, security, cost, service and support, and business value**, with weights reflecting each organization's strategic priorities, and separates "qualification" (pass/fail) criteria from "differentiation" (scored) criteria to keep the process defensible under audit, according to the publisher's own (IntuitionLabs) guidance. This kind of formal Request for Proposal (RFP) discipline matters disproportionately in life sciences because informal vendor selection is much harder to defend during an FDA inspection or an internal quality audit than a documented, weighted evaluation trail.

Model reliability is a second governance axis specific to regulated content. AI systems, particularly large language models, can produce "hallucinations," defined in one peer-reviewed analysis as content generated by extrapolation beyond training data that "seems plausible but is actually a blend of various learned elements" (pmc.ncbi.nlm.nih.gov). Documented instances include an AI system asked for literature on rare metabolic disorders that produced "a thorough paper with several citations with PubMed IDs" where the paper titles were later found to be entirely fabricated and the PubMed identifiers belonged to unrelated papers (pmc.ncbi.nlm.nih.gov). In medical information, regulatory writing, or pharmacovigilance workflows, an AI consultant's proposed mitigation architecture, retrieval-augmented generation with source citation, mandatory human review gates, or confidence scoring, should be an explicit deliverable, not an afterthought raised only when a client asks about it.

A related but broader data-quality problem sits underneath many failed engagements: only an estimated **6% of biopharma data meet FAIR standards** (Findable, Accessible, Interoperable, Reusable), according to Mordor Intelligence's market analysis, and separately, **81% of surveyed firms** cite difficulty reconciling electronic health record (EHR), imaging, and omics data within a single environment (mordorintelligence.com). A

consultant that begins an AI project without a candid data-readiness assessment is effectively skipping the step most likely to determine whether the engagement produces something usable. This data-fragmentation problem is not unique to life sciences: the broader AI consulting services market report from MarketDataForecast notes that **55% of companies report incompatible data systems slow their AI initiatives**, according to Deloitte research cited in that analysis, underscoring that data engineering, not model selection, is usually the binding constraint on any AI consulting timeline (marketdataforecast.com).

To ground these criteria in evaluation practice, procurement checklists commonly recommend a structured process:

- **Define measurable success criteria** before soliciting any vendor, expressed in business outcomes rather than technology features.
- **Request a documented context of use (COU)** for any AI model touching regulated data, mapped against the FDA's risk-based credibility framework.
- **Verify the de-identification method** (Expert Determination or Safe Harbor) used for any training or evaluation dataset containing PHI.
- **Ask for named production references**, not proof-of-concept-only case studies, given the 95% pilot failure rate documented by MIT NANDA researchers.
- **Confirm platform certification status** for any vendor claiming integration expertise with Veeva, a major electronic data capture (EDC) system, or another system of record.
- **Require a hallucination-mitigation architecture** in writing for any Generative AI use case touching regulatory, medical information, or pharmacovigilance content.
- **Score cost against the market benchmark for the engagement type**, distinguishing independent consultant hourly rates ($150 to $500) from MBB-tier billing (roughly $400 associate, $1,000-plus partner) (consultfees.com) (core-mba.pro).
- **Confirm data residency and cross-border transfer mechanisms** if the engagement touches EU patient data subject to the European Health Data Space (EHDS) or GDPR.

Firms that describe themselves as "compliant by design" for FDA 21 CFR Part 11, HIPAA, and GDPR, building validation documentation into delivery from the outset rather than retrofitting it later, illustrate one implementation of this governance-first posture in practice, as the publisher (IntuitionLabs) describes its own approach.

Implementation Considerations and Process Changes

Selecting a consultant is only the first decision; the engagement structure itself determines whether the relationship produces a validated, adopted system. Several process patterns recur across successful and unsuccessful life sciences AI deployments documented in this research.

Scope engagements around prototyping speed, not year-long roadmaps. Pfizer's AWS-backed PACT program compresses proof-of-concept timelines deliberately: a Pfizer team lead noted that building a prototype internally "would take at least 3 months" with appropriate staffing, whereas "a PACT prototype typically has

been no more than 6 weeks" (aws.amazon.com). Of the initiative's first 14 projects, **5 moved into production**, a roughly **36% conversion rate** from prototype to deployed system (aws.amazon.com).

Treat validation and governance as parallel workstreams from day one, not a phase-gate after model development. The ISPE GAMP AI guide frames this explicitly, focusing on "risk-based efforts to allow for efficient, compliant processes" and "improved collaboration among stakeholders and between regulated companies and suppliers" as core benefits of following its methodology (ispe.org). Consultants who present validation as a final compliance sign-off, rather than an integrated design constraint, are structuring the engagement in a way that predictably produces rework. The guide's third stated benefit, "choice of system designs that fit the organization's experience," is a reminder that a rigid, one-size-fits-all validation template is itself a red flag; the risk-based approach is supposed to scale documentation burden to actual system risk, not impose maximal paperwork on every AI use case regardless of stakes (ispe.org).

Insist on measurable, auditable outcomes at each milestone. ZS Associates' survey data shows life sciences technology executives report the most consistent success measuring value in market-sensing product strategy work (**50% report tangible value delivered**), while clinical trials and manufacturing efficiency investments deliver value "less consistently" despite comparably high ambitions (zs.com). This suggests that engagements should sequence toward domains with faster, more measurable feedback loops before tackling harder-to-measure clinical and manufacturing use cases.

Budget for change management and workforce upskilling as a first-class line item. ZS's survey found that 2025 plans emphasize investing in "digital fluency and AI upskilling" (**69%, up from 51% in 2023**) and creating cross-functional teams organized around company goals rather than individual technologies (**72%, up from 56% in 2023**) (zs.com). A consulting proposal silent on organizational change management is, based on this data, addressing a smaller share of the actual adoption problem than the technology work alone suggests.

Match pricing model to engagement risk. Fixed-fee and outcome-based pricing shift delivery risk toward the consultant and away from the client, which is often appropriate for well-scoped, proof-of-concept-stage work; hourly and retainer models suit ambiguous, discovery-phase engagements where scope is expected to evolve. One web-published estimate reports that **73% of clients now prefer outcome-based pricing** across consulting categories generally, though life sciences engagements touching regulatory or validation work often remain hourly or milestone-based given the difficulty of pre-defining "outcome" for a compliance deliverable (consultingdemand.com).

Start with "no-regret bets," not the most ambitious use case on the roadmap. Deloitte frames this explicitly: initiating GenAI projects with "no-regret bets," actions such as automating scientific literature summarization or optimizing contract performance, "offers organizations a prudent path to harnessing AI's potential swiftly and effectively" and creates momentum toward broader, strategic integration (deloitte.com). Deloitte's own recommended sequence, establishing a leadership mandate, aligning on a strategic blueprint of two to three priority areas, identifying no-regret bets, standing up minimum viable governance, and only then launching pilot solutions, mirrors the sequencing implicit in the Pfizer PACT and Novartis-Microsoft cases discussed later in this report (deloitte.com).

Data Analysis and Evidence

The quantitative case for AI consulting investment in life sciences is substantial, but it is uneven across sub-segments, and the data itself sometimes disagrees, which buyers should treat as a signal to verify claims against primary sources rather than accept headline market-size figures uncritically.

On market size, Mordor Intelligence values the global **AI in life sciences market at $3.61 billion in 2025**, growing to **$4.51 billion in 2026** and **$13.64 billion by 2031** at a **24.78% CAGR**, with **North America holding 48.60% of 2025 revenue** and **Asia Pacific projected as the fastest-growing region at a 21.3% CAGR** (mordorintelligence.com). Within that market, software led with a majority share in 2025, while **services (the category that includes consulting) are forecast to grow faster, at a 28.15% CAGR through 2031**, as clients increasingly seek "integration specialists who can align AI outputs with regulated workflows" (mordorintelligence.com). By application, **drug discovery captured 25.60% of 2025 revenue**, and clinical-trials optimization is growing fastest within application segments, at a 27.3% CAGR (mordorintelligence.com). By end user, **pharmaceutical and biotechnology firms controlled 45.40% of 2025 demand**, while contract research organizations (CROs) represent the fastest-expanding customer segment at a **27.2% CAGR through 2031**, as sponsors increasingly outsource analytics-heavy tasks to partners holding multi-sponsor data troves (mordorintelligence.com).

It bears noting that market-size estimates for adjacent categories vary by research firm methodology and scope, a discrepancy worth presenting transparently rather than collapsing into one headline number. *Table 2* below compares four market-size estimates for AI in life sciences and the narrower AI-in-drug-discovery category, gathered from four separate research firms during this report's preparation.

Research Firm	Market Scope	2025 Estimate	Forecast Year and Value	CAGR
Mordor Intelligence	AI in life sciences (broad)	$3.61 billion	$13.64 billion by 2031	24.78% (mordorintelligence.com)
IMARC Group	AI in life sciences (broad)	$3.5 billion	$18.8 billion by 2034	19.84% (imarcgroup.com)
Grand View Research	AI in drug discovery (narrow)	$2.3 billion	$13.8 billion by 2033	24.8% (grandviewresearch.com)
MarketDataForecast	AI consulting services (all industries)	$22.27 billion	$349.80 billion by 2034	35.8% (marketdataforecast.com)

These figures are directionally consistent (low single-digit billions in 2025 for life-sciences-specific AI, growing at roughly 20% to 25% CAGR) but numerically distinct even among firms measuring ostensibly the same category, and North America dominates every regional breakdown at roughly half of global revenue. The gap between the narrow drug-discovery figures and Mordor Intelligence's and IMARC's broader life sciences totals reflects scope choices, not disagreement about the underlying trend; the takeaway for a buyer is that any consultant citing a single market-size figure without naming the research firm and defining scope should be pressed for the source, since "the AI in life sciences market" can mean a $2.3 billion category or a $22 billion one depending on what is being measured.

The broader **AI consulting services market**, spanning all industries, was valued by MarketDataForecast at **$22.27 billion in 2025**, projected to reach **$349.80 billion by 2034** at a **35.8% CAGR** (marketdataforecast.com). Within that market, **North America held a 40.1% share in 2024**, and **digital strategy and transformation services led by service type, capturing 35.3% of global market share in 2024** (marketdataforecast.com).

On adoption and impact, NVIDIA's survey of 600-plus healthcare and life sciences professionals found **83% agree AI will "revolutionize healthcare and life sciences" within three to five years**, **73% say AI is helping reduce operational costs**, and among pharmaceutical and biotech respondents specifically, **62% cite drug discovery as their top Generative AI use case** (blogs.nvidia.com). The same survey found spending priorities concentrated on identifying additional use cases, optimizing workflow, and hiring more AI talent, ahead of infrastructure alone, with **78% of respondents intending to increase their AI infrastructure budget** and more than a third planning increases greater than 10% (blogs.nvidia.com).

Clinical trial economics illustrate both the opportunity and the structural difficulty AI consulting engagements must address. Industry data show **30% to 50% of Phase III trials miss their enrollment timelines**, and a broader dataset cited by digital-recruitment analysts finds **85% of studies fail to hit enrollment goals**, **80% run past planned timelines**, and **only 2% of the eligible population ever joins a trial** (lifebit.ai). Against this backdrop, drug development economics remain unforgiving regardless of how a candidate was discovered: **Phase 1 trials may cost only a few million dollars** and involve **15 to 20 people**, while **Phase 2 can cost around $45 million**, and **Phase 3 efficacy studies take three to four years, cost hundreds of millions of dollars, and succeed only about 55% of the time** (chemistryworld.com). Recursion Pharmaceuticals' stated aim is to reduce the roughly **90% failure rate** that characterizes traditional drug discovery overall (chemistryworld.com). No AI consultant, however capable, changes the underlying biology; the honest framing is that AI can compress discovery timelines and improve target selection, not eliminate clinical-stage risk.

Finally, on the consulting-specific failure risk that motivates rigorous vendor evaluation in the first place, the MIT NANDA "GenAI Divide" report, based on **150 leadership interviews, a 350-employee survey, and analysis of 300 public AI deployments**, found that while about **5% of AI pilot programs achieve rapid revenue acceleration**, the "vast majority stall, delivering little to no measurable impact on P&L" (fortune.com). The report's lead author identified the core issue as a "learning gap" in enterprise integration rather than model quality, and found that **more than half of Generative AI budgets are devoted to sales and marketing tools**, while the biggest realized return on investment (ROI) actually comes from back-office automation (fortune.com). This is directly relevant to life sciences buyers evaluating consulting proposals: a proposal weighted heavily toward customer-facing GenAI features, at the expense of back-office regulatory, quality, or data-engineering automation, may be optimizing for the wrong side of the documented ROI distribution. The same MIT report notes that some AI-native startups have seen "revenues jump from zero to $20 million in a year" by picking one pain point, executing well, and partnering smartly with the companies that use their tools, a contrast the report's lead author draws deliberately against the median enterprise's diffuse, feature-scattered GenAI rollout (fortune.com).

Beyond pilot-stage risk, organizational readiness gaps are well documented within life sciences specifically. ZS Associates' executive survey found that AI and emerging-technology opportunities are already reshaping data strategies (**77% of respondents**), revenue targets (**74%**), productivity targets (**73%**), and ways of working

(**72%**) (zs.com). On the technology-organization side, most respondents plan to improve how their function demonstrates business outcome measures (**60%**) and shift toward more proactive innovation (**59%**) in 2025, evidence that measurement discipline, not just AI capability, is the area life sciences technology leaders themselves identify as lagging (zs.com).

Case Studies and Real-World Examples

Pfizer and Amazon Web Services: The PACT Model for Rapid, Governed Prototyping

Pfizer's collaboration with AWS through the **Pfizer-Amazon Collaboration Team (PACT)** illustrates a consulting-adjacent model built around speed and measured outcomes rather than open-ended strategy work. Under PACT, Pfizer pursued **14 projects** using Generative AI and Machine Learning, with the stated goals of saving scientists up to **16,000 hours of search time annually** and cutting certain infrastructure costs by **55%** (aws.amazon.com). The first use case targeted a data-discovery problem specific to pharmaceutical R&D: a single drug's development can generate approximately **20,000 documents**, and scientists historically searched manually across disparate tools; the AI-assisted solution aimed to cut discovery time by **up to 80% for 1,500 affected scientists** (aws.amazon.com). Separately, PACT-developed anomaly detection for Pfizer's continuous manufacturing process (PCMM) used Amazon SageMaker and related tooling to predict maintenance needs and reduce unplanned downtime. Of the initial 14 projects, **5 reached production**, and Pfizer's team lead credited the arrangement with changing organizational risk appetite: "We can learn fast and fail fast with these prototyping experiments, and that's a major benefit for us," per the AWS case study (aws.amazon.com).

Novartis and Microsoft: A Five-Year Enterprise AI Alliance

In 2019, **Novartis** and **Microsoft** announced a five-year AI alliance intended to help Novartis make sense of data generated across laboratory experiments, clinical trials, and manufacturing operations, with early application to personalizing treatment for age-related macular degeneration and improving manufacturing of chimeric antigen receptor T-cell (CAR-T) cancer therapies (biopharmadive.com). Novartis had, at that point, accumulated an estimated **2 million patient-years of clinical trial data** over two decades and data covering **1.5 million chemical compounds** (biopharmadive.com). No specific financial terms were disclosed, but the alliance also created a company-wide AI innovation lab intended to make AI-powered applications accessible across Novartis's more than **100,000 employees** (biopharmadive.com). This case demonstrates the multi-year, enterprise-wide commitment that hyperscaler-anchored AI transformation typically requires, in contrast to the faster prototype cadence of the Pfizer PACT model.

Insilico Medicine: The First AI-Designed Drug to Approach Phase 3

Insilico Medicine's **rentosertib** (also referenced by its development code INS018_055), an AI-discovered and AI-designed small molecule for idiopathic pulmonary fibrosis (IPF), reported a **71-patient study in China** in June 2025 and has since entered Phase 3: Insilico dosed the first patient in the GENESIS-IPF-3 trial in

September 2026 (insilico.com). As of Chemistry World's July 2025 report, no AI-designed drug had ever entered a phase 3 clinical trial (chemistryworld.com). Independent commentators quoted in the same report remain divided: medicinal chemist Derek Lowe notes that AI is best suited to stages that are "in almost inverse proportion to how important they are to the expense of developing drugs," meaning it accelerates early discovery but does little for the costliest, highest-failure-risk late-stage trials (chemistryworld.com). Separately, deals between AI-native biotechs and large pharmaceutical partners have proliferated: **Isomorphic Labs**, an Alphabet subsidiary, entered collaborations with Eli Lilly and Novartis in January 2024 potentially worth billions, and **BenevolentAI** signed a **$594 million** strategic collaboration with Merck in 2023 (chemistryworld.com). **Recursion Pharmaceuticals**, founded in 2013 on a machine-learning strategy built around a single foundational assay, now runs highly automated in-house laboratories and states its aim is to reduce the roughly 90% failure rate of traditional drug discovery; the company dosed its first patient in a Phase 1/2 clinical study of **REC-1245**, a molecular glue targeting a protein essential to the survival of certain cancer cells, in December 2024 (ir.recursion.com) (chemistryworld.com).

IntuitionLabs and Scilex Holding Company: Commercial Analytics via Certified Veeva Implementation

A further, smaller-scale case illustrates how platform-native implementation partners apply AI-adjacent analytics inside an existing system of record rather than as a stand-alone tool. **Scilex Holding Company** engaged **IntuitionLabs**, a Veeva Vault CRM X-Pages Partner, to address fragmented systems across prescription, claims, and customer relationship management (CRM) data, delayed weekly or monthly performance insight cycles, and limited pre-call planning intelligence (intuitionlabs.ai). The delivered solution, custom Veeva MyInsights dashboards powered by Veeva Nitro, unified prescription, claims, and CRM data into daily automated updates and added territory-level performance dashboards with goal tracking. Scilex's Senior Director of Commercial Operations, Preetaman Wadhwa, is quoted describing the result: "MyInsights represent a true breakthrough for our commercial operations. For the first time in our company history, we have a solution that aggregates claims data, prescription information, goal tracking, and CRM activities in one place with daily automated updates" (intuitionlabs.ai). Disclosure: IntuitionLabs is the publisher of this article, this case is drawn from its own materials and has not been independently verified, and the project concerned commercial analytics dashboards rather than AI. It is included only as an illustration of the platform-partner delivery model described earlier in this report, not as independent evidence, and is distinct from the Novartis-Microsoft and Pfizer-AWS cases.

Implications and Future Directions

The trajectory of AI consulting for life sciences over the next 24 to 36 months will be shaped by three converging forces. First, **regulatory deadlines will compress the window for informal AI adoption**. With EU AI Act high-risk obligations phasing in through **December 2027 and August 2028**, and the FDA's AI credibility framework already shaping submission strategy since January 2025, organizations that have not yet formalized AI governance will face increasing pressure to retrofit compliance into systems built without it, a materially more expensive path than building compliance in from the outset (digital-strategy.ec.europa.eu).

Second, **the consolidation of consulting supply around genuine domain specialization is likely to accelerate**, driven by buyer experience with the MIT-documented 95% pilot failure rate. As life sciences buyers become more sophisticated procurers, evidenced by the shift toward formal weighted scorecards and RFP discipline, generalist AI consultancies without demonstrable regulatory or GxP validation experience should expect increasing difficulty winning production-scale (as opposed to pilot-scale) engagements. The services segment of the AI in life sciences market growing faster than software through 2031, noted earlier in this report, is consistent with buyers increasingly paying for integration and compliance expertise specifically, not raw AI capability that is now widely commoditized. Consulting firms that formalize this specialization, whether through named regulatory practice groups, published validation methodologies, or platform certifications, are positioned to capture a disproportionate share of that growth relative to firms marketing generic "AI transformation" services.

Third, **federated and privacy-preserving data infrastructure will reshape what consultants can credibly promise**. The European Health Data Space (EHDS) Regulation entered into force on 26 March 2025, beginning a transition phase towards application, with rules on secondary use of health data due to start applying for most data categories later in that transition, a policy shift running in parallel with the EU AI Act's own phased implementation described earlier in this report (health.ec.europa.eu). Consultants able to architect solutions around federated training, rather than requiring centralized data pooling, will hold a structural advantage as this infrastructure matures, particularly given that GDPR and national data-protection frameworks already constrain how patient-level data can move across borders even within a single sponsor organization.

For life sciences organizations evaluating AI advisory partners now, the practical implication is to weight selection criteria toward evidence of production deployment, regulatory fluency specific to the FDA's COU framework and the EU AI Act's staggered timeline, and platform certification (where applicable) over breadth of AI capability claims alone. Adjacent advisory firms embedded in existing enterprise systems, rather than positioned as a competing product, argue their proximity to production data and workflows is itself a risk-reducing feature of the engagement model, a claim buyers should verify against named, referenceable production outcomes rather than accept on its face, exactly as this report recommends doing with any consulting claim in this category.

Frequently Asked Questions (FAQs)

What does "AI consulting for life sciences" actually cover? It spans strategy advisory (where and how to invest in AI), technical implementation (building or integrating models into existing systems like Veeva, electronic data capture platforms, or laboratory information management systems), regulatory and validation support (mapping AI systems to FDA, EMA, and GAMP 5 requirements), and change management (workforce training and process redesign). Few firms are equally strong across all four; buyers should map their actual need before shortlisting.

How is life sciences AI consulting different from general enterprise AI consulting? The core difference is regulatory exposure. A retail AI project that underperforms is a business problem; a life sciences AI project that mishandles GxP-relevant data, patient data, or a regulatory submission can trigger findings from the FDA or

EMA and, in the worst case, patient-safety consequences. This is why frameworks like the FDA's risk-based credibility assessment and the ISPE GAMP AI guide exist specifically for this sector (fda.gov) (ispe.org).

How do I evaluate an AI vendor or consultant for pharma and biotech specifically? Use a weighted scorecard covering technical capability, regulatory compliance, data handling, security, cost, service and support, and business value, with pass/fail qualification criteria separated from scored differentiation criteria, per the publisher's own (IntuitionLabs) guidance. Demand named, production-scale references, not pilot-only case studies, given the documented 95% GenAI pilot failure rate across industries generally (fortune.com).

How much does AI consulting cost in life sciences? Indicative web estimates that are not specific to life sciences AI suggest independent strategy consultants typically charge **$150 to $500 per hour**, while global systems integrators and MBB-tier firms bill closer to **$400 per hour for associates and over $1,000 per hour for partners**, figures that include firm overhead (consultfees.com) (core-mba.pro). Actual project cost depends heavily on scope, data readiness, and whether validation and compliance documentation is included from the outset or added later.

Can AI actually shorten drug development timelines, or is that overstated? Evidence is mixed by phase. AI has demonstrably compressed early discovery: Insilico Medicine took its anti-fibrotic candidate ISM001_055 "from start to Phase 1 in 30 months," a timeline the company describes as "several orders of magnitude faster and cheaper as compared to a traditional drug discovery process" that typically costs around $430 million out-of-pocket and takes three to six years (insilico.com). But Phase 2 and Phase 3 trial economics, cost, duration, and success probability, remain governed by clinical biology, not software, and Phase 3 still succeeds only about **55% of the time** regardless of how the candidate was discovered (chemistryworld.com).

What is the single biggest reason life sciences AI engagements fail? A commonly cited failure mode in the MIT NANDA pilot-failure data is a "learning gap" in how tools and organizations integrate AI, rather than the quality of the underlying models (fortune.com).

Does the EU AI Act apply to a U.S.-headquartered pharmaceutical company? Yes, if the company places AI systems on the EU market or those systems' outputs are used within the EU, regardless of where the company is headquartered. The Act's high-risk obligations phase in on a staggered schedule, with systems used in areas like biometrics and critical infrastructure applying from **2 December 2027**, and systems embedded into regulated products, the category most relevant to AI-enabled medical devices and diagnostics, applying from **2 August 2028** (digital-strategy.ec.europa.eu). Transparency obligations for AI-generated content labeling take effect earlier than the high-risk rules: under the adopted AI Omnibus (Regulation (EU) 2026/1744), providers of generative AI systems placed on the market before 2 August 2026 have until **2 December 2026** to comply with the Article 50(2) machine-readable marking obligation (eur-lex.europa.eu).

What consulting archetype is best for a mid-size biotech with limited internal AI expertise? Evidence in this report points toward boutique life sciences specialists or platform-native implementation partners over global systems integrators for organizations without a large internal technology function, since the former typically staff engagements with fewer, more senior, domain-specific practitioners rather than a large team billed at MBB-tier rates (core-mba.pro). A mid-size organization already running Veeva, for example, may get faster time-to-value from a certified X-Pages implementation partner than from a strategy-only engagement that produces a roadmap but no deployed system (veeva.com).

Conclusion

Evaluating AI consulting for life sciences requires holding two standards simultaneously: the general standard for judging any AI consultant, evidence of production deployment, transparent pricing, and measurable outcomes, and a sector-specific standard built on regulatory fluency, GxP validation methodology, and rigorous data governance that most generalist AI consultancies simply do not carry. The market context supports urgency: a life sciences AI market growing at roughly a quarter every year through 2031, life sciences executives reporting rising investment in data, digital, and AI, and projected benefits reported by early movers such as Pfizer. But the same context counsels caution: a documented 95% pilot failure rate industry-wide and a regulatory landscape, spanning the FDA's 2025 credibility framework, the EMA's reflection paper, and the EU AI Act's staggered 2026 through 2028 timeline, that continues to evolve in ways that penalize consultants who cannot speak fluently to compliance from day one.

The practical takeaway for a life sciences organization building a shortlist is to weight the evaluation toward verifiable evidence over marketing claims: named production references rather than pilot anecdotes, explicit de-identification methodology rather than a generic compliance assurance, documented context-of-use mapping rather than a vague reference to "FDA readiness," and platform certification where the engagement touches an existing system of record. Firms across every archetype examined here, global integrators, boutique specialists, and platform-native implementation partners, can meet that bar; the burden is on the buyer to demand the evidence rather than assume it. Given the pace of regulatory change documented throughout this report, that evaluation discipline is not a one-time procurement exercise but a standing requirement that should be revisited at least annually as FDA, EMA, and EU AI Act obligations continue to phase in through 2028.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.mordorintelligence.com/industry-reports/artificial-intelligence-in-life-sciences-market
2. https://www.marketdataforecast.com/market-reports/ai-consulting-services-market
3. https://www2.deloitte.com/content/dam/Deloitte/us/Documents/us-realizing-transformative-value-from-AI-GenAI-in-life-sciences-032124.pdf
4. https://www.deloitte.com/us/en/industries/life-sciences-health-care/articles/value-of-genai-in-pharma.html
5. https://blogs.nvidia.com/blog/ai-healthcare-life-sciences-survey-2025
6. https://aws.amazon.com/solutions/case-studies/pfizer-PACT-case-study
7. https://insilico.com/news/isn1009261-insilico-medicine-doses-first-patient-genesis-ipf-3
8. https://www.chemistryworld.com/news/as-ai-designed-drug-looks-to-pass-final-hurdle-will-this-tech-change-drug-discovery-forever/4021894.article
9. https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo
10. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
11. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

12. https://ispe.org/publications/guidance-documents/gamp-guide-artificial-intelligence
13. https://www.zs.com/insights/2025-survey-data-digital-ai
14. https://newsroom.accenture.com/content/4q-full-fy25-earnings/accenture-reports-fourth-quarter-and-full-year-fiscal-2025-results.pdf
15. https://www.accenture.com/us-en/industries/life-sciences
16. https://www.lek.com/industries/life-sciences-pharma
17. https://clarkstonconsulting.com/industries/life-sciences
18. https://www.eversana.com/our-expertise/ai-accelerator
19. https://www.accenture.com/us-en/careers/jobdetails?id=R00315241_en
20. https://www.core-mba.pro/tool-hub/consultant-rate
21. https://www.statsby.ai/consulting.html
22. https://www.sciagen.ai/life-sciences
23. https://linkedin.com/company/phintel
24. https://www.iuvotech.com/ai-consulting/pharma
25. https://viosenza.com
26. https://siglasciences.ai
27. https://www.veeva.com/meet-veeva/partners/x-pages
28. https://www.hhs.gov/hipaa/for-professionals/special-topics/de-identification/index.html
29. https://consultfees.com/use-cases/strategy-consultants
30. https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development
31. https://www.ema.europa.eu/en/news/reflection-paper-use-artificial-intelligence-lifecycle-medicines
32. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML?uri=OJ:L_202601744
33. https://pmc.ncbi.nlm.nih.gov/articles/PMC10552880
34. https://consultingdemand.com/blog/consulting-fees-by-industry
35. https://www.imarcgroup.com/artificial-intelligence-in-life-sciences-market
36. https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-drug-discovery-market
37. https://lifebit.ai/blog/clinical-trial-recruitment-digital-case-study
38. https://www.biopharmadive.com/news/novartis-microsoft-ai-artificial-intelligence-alliance-car-t/564117
39. https://ir.recursion.com/news-releases/news-release-details/recursion-announces-first-patient-dosed-phase-12-clinical-study
40. https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en
41. https://insilico.com/phase1

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.