IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Biotech AI Strategy Timeline: How Long to Build a Roadmap

> 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 · 2026-06-28 · biotech ai strategy timeline · biotech ai roadmap · ai strategy for biotech companies · ai implementation timeline pharma · biotech ai department structure · building an ai team in biotech · life sciences ai consulting · pharma ai roadmap execution

Original article: https://intuitionlabs.ai/articles/biotech-ai-strategy-timeline-roadmap



In this article

1. Executive Summary
2. Introduction and Background
3. Key Changes: The Phases of a Biotech AI Strategy Timeline
4. Phase 1: Diagnostic and Readiness Assessment (Days 1 to 90)
5. Phase 2: Strategy Formulation and Board Approval (Months 3 to 6)
6. Phase 3: Governance, Team Formation, and Pilot Selection (Months 6 to 12)
7. Phase 4: Scaling and Enterprise Integration (Year 2 to Year 3)
8. Implementation Considerations and Process Changes
9. Data Analysis and Evidence
10. Case Studies and Real-World Examples
11. Amgen: A Chief Technology Officer Mandate Built Around "The Hinge Moment"
12. GSK: Six Years of Infrastructure Before the Headlines
13. Moderna: Scaling Without Scaling Headcount
14. Merck: GPTeal and the Economics of Enterprise-Wide Rollout
15. Insilico Medicine and Exscientia: AI-Native Biotechs Compressing Preclinical Timelines
16. AstraZeneca: Embedding AI Across the Full R&D Value Chain Rather Than a Single Function
17. Implications and Future Directions
18. Frequently Asked Questions (FAQs)
19. Conclusion

Executive Summary

Biotech and pharmaceutical executives asking "how long does it take" to build an artificial intelligence (AI) strategy are usually asking the wrong question first. The evidence gathered across regulatory filings, consultancy benchmarks, and named corporate deployments points to a specific answer only once "strategy" and "roadmap" are separated: the author's planning heuristic, not a measured industry benchmark, is a diagnostic phase of roughly 60 to 90 days, a strategy formulation and governance phase of three to six months, and an enterprise scaling phase that brings the cumulative total to roughly two to three years before AI moves from pilot activity to a durable operating capability (bcg.com) (intuitionlabs.ai). That timeline is not a formality: Deloitte's fourth annual Life Sciences Outlook Survey of **280** C-suite executives across biopharma and medtech found that only **22%** of life sciences leaders said they had successfully scaled AI as of late 2025, and just **9%** reported achieving significant returns on those efforts (deloitte.com).

The industry has largely agreed on the destination even where the timeline remains contested. Capgemini's Research Institute found that **82%** of biopharma executives believe AI will fundamentally transform research and development (R&D), and **63%** anticipate that most new molecular entities will originate from AI-driven platforms within the next decade (capgemini.com). Strategy& (PwC) has estimated that fully industrializing AI across pharma could unlock on the order of **$250 to $254 billion** in annual value by 2030, and Bessemer

Venture Partners has separately sized the addressable opportunity at over **$100 billion** in net annual value unlock against a **$150 billion**-plus outsourced pharma services market, while noting that only about **5%** of pharma companies have captured measurable AI value to date (bvp.com).

The roadmap that closes the gap between ambition and realized value follows a recognizable sequence regardless of company size: a diagnostic and AI inventory (roughly the first 90 days), a board-approved strategy document with financial and risk framing (months three through six), governance and pilot selection built around a small number of enterprise-level use cases (months six through twelve), and a multi-year scaling program that builds a Center of Excellence (CoE) or hub-and-spoke operating model spanning data, talent, and compliance, consistent with Capgemini's finding that **79%** of biopharma organizations are now actively developing formal strategies of this kind rather than relying on ad hoc pilots (capgemini.com). Real-world timelines corroborate this arc: Sanofi's Toronto AI hub grew from **2** people to **150** in roughly two years en route to becoming an 80-person global strategy function (jobs.sanofi.com), Amgen's 2023 executive reorganization around AI is still, in its own leadership's telling, an unfinished "continuous, radical transformation" (endpoints.news), and GSK describes **six years** of quiet infrastructure building before its AI-driven R&D model began generating headline results (govinfosecurity.com).

Regulatory timing constrains the roadmap as much as internal readiness does. The U.S. Food and Drug Administration (FDA) published its first draft guidance on AI in drug and biological product regulatory submissions in January 2025, built on more than **500** AI-related submissions reviewed between 2016 and 2023 and over **800** public comments (fda.gov), while the European Union's AI Act, as amended by the Digital Omnibus on AI (Regulation (EU) 2026/1744), now applies its high-risk obligations for Annex III systems from 2 December 2027 and for Annex I product-embedded systems from 2 August 2028 (eur-lex.europa.eu). Whether a given biotech AI system is high-risk at all depends on its intended use, so a roadmap should include an early classification step.

Case evidence from named companies illustrates both the ceiling and the floor of what a well-executed timeline can deliver. Insilico Medicine took a novel drug candidate from target identification to preclinical candidate in under 18 months at a fraction of typical cost, compared with the traditional four-to-six-year cycle (pmc.ncbi.nlm.nih.gov), and Eli Lilly has stood up what it calls the largest AI factory wholly owned by a pharmaceutical company, built on more than **1,000** NVIDIA Blackwell Ultra graphics processing units (GPUs) (blogs.nvidia.com). At the same time, some industry commentary warns that generative AI pilots in pharma risk stalling without deliberate management, a caution echoed by Gartner's finding that **73%** of chief information officers (CIOs) in Europe, the Middle East, and Africa report their organizations are breaking even or losing money on AI investments (gartner.com). This report lays out the phased timeline, the organizational structures that make it executable, the data supporting each stage, and the named case studies that show what disciplined execution looks like in a highly regulated, science-driven industry.

Introduction and Background

Every biotech and pharmaceutical company building an artificial intelligence program eventually confronts the same question from its board: how long will this actually take? The honest answer is that it depends on what "this" refers to. Industry practice increasingly distinguishes between an **AI strategy**, the document that identifies where AI will create durable competitive advantage and how the organization will be structured to

capture it, and an **AI roadmap**, the sequenced, resourced implementation plan that turns that strategy into deployed capability. Conflating the two is a common planning error: a board can approve a strategy without anyone committing the sequenced, resourced implementation plan needed to deliver it.

The urgency behind getting this timeline right is not abstract. Industry surveys report that approximately **90%** of pharma and biotech technology leaders see AI and digital pressures as direct threats to business growth, and in response **77% to 82%** of executives now view digital innovation as a competitive differentiator. Accelerating drug discovery is the top AI priority for **52%** of firms, with enhancing patient engagement a close second at **43%**. Separately, as detailed further below, a survey of the top 20 pharmaceutical companies found that 70% of leaders call AI an "immediate priority," rising to 85% among the very largest firms, and more than 80% report that they are increasing AI budgets rather than holding them flat. Yet urgency and readiness are not the same thing. The same body of research found that **under 60%** of life sciences organizations have AI-specific standard operating procedures (SOPs) or conduct regular AI audits, even as **75% to 86%** of firms report plans to implement AI capabilities in the near term (intuitionlabs.ai). About 80% of leaders in the same Define Ventures survey reported already having a dedicated AI governance structure (fiercepharma.com), a formal-governance rate that outpaces the audit rigor noted above.

Biotech is structurally different from the software and financial services sectors where most generic "AI transformation timeline" guidance originates. Drug candidates face multi-year clinical development cycles; data is generated under Good Clinical Practice (GCP), Good Laboratory Practice (GLP), and Good Manufacturing Practice (GMP) quality regimes; and AI used to produce information or data intended to support regulatory decision-making on drug safety, effectiveness, or quality is within the scope of the FDA's January 2025 nonbinding draft guidance on AI credibility assessment, which recommends a risk-based framework tied to a model's specific "context of use" (fda.gov). This means a biotech AI timeline cannot simply borrow a generic enterprise software rollout plan; it has to build in validation, audit trail, and human-oversight requirements from the earliest planning stages rather than retrofitting them later. That January 2025 draft guidance did not emerge from a vacuum: its content was shaped by feedback from a December 2022 expert workshop convened by the Duke Margolis Institute for Health Policy, more than **800** public comments submitted on a May 2023 discussion paper, and CDER's own experience reviewing over 500 AI-related submissions between 2016 and 2023, followed by additional hybrid public workshops in August 2024 and October 2025 (fda.gov). Boston Consulting Group (BCG) has catalogued more than **130** potential generative AI use cases across the biopharma value chain, spanning research, clinical development, manufacturing, and commercial operations, which itself signals why sequencing and prioritization, not simply technology selection, is the central strategic problem (bcg.com). European regulators have moved in parallel with the FDA, not independently: the European Medicines Agency (EMA) published its own draft reflection paper on AI across the medicinal product lifecycle, emphasizing a human-centric approach and early regulatory scientific advice whenever an AI system could affect a medicine's benefit-risk balance (ema.europa.eu), reinforcing that a Phase 2 strategy document needs to anticipate at least two major regulatory frameworks, not just the FDA's (ema.europa.eu).

This report addresses the timeline question directly: how long each phase of a biotech AI strategy realistically takes, what organizational structures let a company execute that timeline rather than stall in pilot purgatory, what the data says about returns and risk, and what named companies, from megacap pharma to venture-backed AI-native biotechs, have actually done and how long it took them.

Key Changes: The Phases of a Biotech AI Strategy Timeline

Phase 1: Diagnostic and Readiness Assessment (Days 1 to 90)

The first phase of any credible biotech AI timeline is a structured inventory of what AI already exists in the organization, formally sanctioned or not, and an honest assessment of data, talent, and governance readiness. In practice this phase runs roughly 60 to 90 days, mirroring the "first 90 days" framing that has become common for newly appointed Chief AI Officers (CAIOs), a role increasingly common at large pharmaceutical companies as they centralize AI accountability under a single executive. The core activities are an AI system inventory (production deployments, vendor-embedded features inside laboratory information management systems (LIMS) and clinical trial management systems (CTMS), and unsanctioned employee use of public large language models), a stakeholder map, and a gap analysis against data quality, model validation, and governance requirements.

This diagnostic phase matters because the inventory is almost always larger and messier than leadership expects. A 2024 survey found that over **80%** of life sciences companies that had implemented AI had no formal policies or audit processes in place at the time, even as adoption of individual AI use cases was already widespread (intuitionlabs.ai). Separately, **65%** of the top 20 Big Pharmas reported having banned consumer generative AI tools such as ChatGPT over concerns that sensitive internal data could be leaked, even while quietly running other AI pilots elsewhere in the business, underscoring how fragmented AI governance typically is at the point a formal strategy process begins (fiercepharma.com). A companion 2024 ZoomRx survey of 200-plus industry professionals found **83%** still labeled AI as "overrated" even as more than half had AI already in production, a skepticism-versus-adoption split that a diagnostic phase is designed to surface (fiercepharma.com). Organizations are also closing capability gaps through organizational change, consistent with the industry forecasts of halved early-stage development timelines reported by Reuters (reuters.com) and with Amgen's own decision to reorganize around a newly created chief technology officer role (endpoints.news).

Gartner's guidance to CIOs generally reinforces the value of an honest diagnostic before committing capital: the firm's July 2025 survey of over **700** CIOs found that, in EMEA, **73%** report their organizations are breaking even or losing money on AI investments, and Gartner explicitly recommends that for every AI tool an organization buys, it should anticipate roughly **10** hidden costs on top of licensing, including the transition costs of training and change management (gartner.com). A diagnostic phase that skips this cost accounting tends to produce a strategy document that cannot survive contact with a finance committee.

Phase 2: Strategy Formulation and Board Approval (Months 3 to 6)

The second phase converts the diagnostic into a strategy document built for capital allocation rather than technology enthusiasm. Boards approve AI strategies the way they approve any large capital request: they want to know what is being bought, what it costs, how it pays back, and what could go wrong. A strategy that lacks financial rigor relative to other capital requests of comparable size risks conditional approval, deferral, or

a rework request on first pass, a point consistent with Capgemini's broader observation that formal AI strategy development is now the norm rather than the exception among biopharma organizations (capgemini.com). Gabriele Ricci, chief data and technology officer at Takeda Pharmaceuticals, captured the tension boards are wrestling with when he told Deloitte that "volatility fuels innovation," adding that the industry is "entering a period of purposeful transformation, where discipline and innovation must coexist as the industry matures beyond hype toward measurable productivity from AI and data" (deloitte.com), a framing that maps directly onto why Phase 2 strategy documents need financial discipline rather than pure technology narrative to clear board review.

This phase typically runs three to six months because it requires genuine cross-functional input: a use-case prioritization exercise informed by a value-versus-complexity matrix, a governance framework spanning IT, R&D, legal, and quality assurance, and a financial model that treats AI as a capital allocation decision rather than an information technology (IT) line item. McKinsey's "six enablers" framework for pharma AI transformation, widely cited in industry practice, starts from the same question boards ask: what problem is being solved, and what is the return on investment (ROI) (intuitionlabs.ai). BCG's guidance is similarly explicit about scope discipline during this phase: companies should elevate only **three to five** use cases to enterprise-level funding and focus, since proliferating pilots and proofs of concept across the business dilutes attention without building traction (bcg.com). Capgemini frames the prize this phase is meant to unlock in end-to-end terms: by streamlining discovery, optimizing trial design, and enabling predictive insights, AI is positioned to make biopharma R&D more agile, data-driven, and outcome-oriented across the whole value chain rather than in a single isolated function (capgemini.com).

Capgemini's Research Institute found that **79%** of biopharma organizations are now actively developing formal strategies to integrate AI across their R&D value chain, and that **74%** of executives specifically believe generative AI holds significant potential in early-stage drug discovery, evidence that this strategy-formulation phase has become close to universal practice rather than an outlier behavior among digital leaders (capgemini.com). The financial stakes attached to getting this document right are large: PwC's Strategy& has projected that AI could lift pharma operating margins from roughly **20%** to **40%**, translating into an estimated **$254 billion** a year in additional profit globally through efficiency gains and new revenue streams once AI is industrialized at scale.

Phase 3: Governance, Team Formation, and Pilot Selection (Months 6 to 12)

Once a strategy has board sign-off, the roadmap phase begins in earnest. This is where organizational structure decisions, centralized AI Center of Excellence (CoE), fully decentralized embedded teams, or a hybrid hub-and-spoke model, get made, and where the first funded pilots launch. ZS's benchmarking study of **15** large and midsize biopharma companies found that a hybrid structure is becoming the prevalent model: a central data science team acts as a hub that sets data strategy, trains data scientists, and builds shared AI capabilities, while domain-specific "spoke" teams embedded in R&D functions apply those capabilities to their own custom analytics problems (zs.com). Many information technology (IT) departments have correspondingly established their own AI centers of excellence to bring data scientists, data engineers, and other specialists

together to scale AI infrastructure across the business, a structural echo of the IntuitionLabs organizational design framework that also treats the CoE as the entity responsible for standardizing best practices and providing oversight for data quality and regulatory compliance enterprise-wide (intuitionlabs.ai).

Talent strategy during this phase tends to follow a "T-model": generalists who can bridge data science and business functions, paired with deep specialists such as biostatisticians and computer vision experts who combine advanced technical training with functional domain knowledge (zs.com). Some organizations accelerate this build-out through acquisition or partnership rather than organic hiring alone, a pattern the peer-reviewed literature confirms is now closer to the norm than the exception: **97%** of the AI drug-discovery studies reviewed in the *Pharmaceuticals* systematic review disclosed a formal industry partnership rather than purely independent development (pmc.ncbi.nlm.nih.gov). Sanofi's experience shows how quickly a dedicated AI organization can scale once a mandate and location are set: its Toronto AI Center of Excellence grew from **2** people to **150** in roughly two years, chosen deliberately for its concentration of AI research talent tied to the University of Toronto and the Vector Institute (jobs.sanofi.com). The team's founding member was promoted from Global Director to Global Vice President within a year, and by December 2023 held a Global Senior Vice President role overseeing roughly **80** people across Sanofi's global digital strategy function (jobs.sanofi.com). Notably, Sanofi's central Toronto hub does not attempt to own every AI initiative itself; it works in close partnership with in-house digital accelerator teams based in Paris and Lyon, France, that build breakthrough digital solutions for specific therapeutic areas while Toronto focuses on core AI infrastructure and scalable products, a distributed pattern consistent with the hub-and-spoke model described later in this report (jobs.sanofi.com).

Governance built during this phase needs to be substantive rather than symbolic. BCG recommends creating a dedicated AI ethics office or similar function to oversee responsible AI policy focused on safety without becoming an innovation bottleneck, alongside a technology platform flexible enough to accommodate different foundational models as they evolve, since prompts and fine-tuning approaches are not portable across models without adaptation (bcg.com). Pilot selection during this phase should prioritize **five to ten** initial projects that balance near-term return on investment against longer-term strategic bets, documented with clear owners and milestones as part of the broader roadmap. Organizations still need to mature their data and technology foundations alongside governance during this phase: leading firms extend their data networks and tear down internal silos between discovery, clinical development, and medical affairs while also establishing external real-world data partnerships, a pattern consistent with AstraZeneca's own description of federated electronic health record technology unlocking new opportunities to enhance clinical research once internal data governance is in place (astrazeneca.com).

Phase 4: Scaling and Enterprise Integration (Year 2 to Year 3)

The final and longest phase is where most biotech AI programs actually fail, not in strategy formulation or initial piloting, but in the transition from isolated proof of concept to enterprise-scale, repeatable deployment. ZS's benchmarking research found that despite widespread enthusiasm for AI in biopharma R&D, only **20%** of companies studied had successfully scaled AI initiatives and driven mainstream adoption across their organizations (zs.com). Industry commentary also warns that generative AI pilots can stall absent careful management of the transition from pilot to production.

A structured three-year scaling arc is a common planning horizon in this phase, typically sequenced with laboratory automation and systems integration in year one, expansion into clinical trial management and drug discovery workflows in year two, and predictive analytics applied to operational cost reduction in year three. Deloitte's 2026 Life Sciences Outlook, surveying **280** C-suite executives from biopharma and medtech companies across the United States, Europe, and Asia, found that **48%** of respondents identified accelerated digital transformation as a trend expected to substantially affect their organizations in 2026, a statistically significant increase from 2025, and that **41%** specifically cited the proliferation of generative AI as an influential trend, with **30%** newly citing agentic AI, systems capable of acting autonomously to achieve goals, as an emerging area of interest (deloitte.com). Despite that investment, the same survey found only 22% of life sciences leaders reported successfully scaling AI and just 9% reported significant returns, as cited above, the clearest available evidence that the scaling phase, not the strategy or pilot phases, is where most biotech AI timelines actually break down. The same survey found that scaling is competing for executive attention against several other pressures that any realistic Phase 4 timeline needs to account for: **39%** of respondents cited geopolitical tensions as a top issue, a 20 percentage point jump from the prior year and the largest increase among all tracked trends, while **35%** cited rising cybersecurity concerns and **36%** of U.S. respondents specifically flagged agency restructuring at the FDA and the Department of Health and Human Services as a factor likely to shape strategy in 2026 (deloitte.com). None of these pressures make the four-phase timeline shorter; they make the discipline of sticking to it, rather than constantly re-litigating scope in response to the news cycle, more important. Talent competition adds a further constraint on Phase 4 speed: one medtech supply chain executive summarized it as a shortage of skilled staff, not AI tools, noting "there aren't enough phlebotomists or lab managers" (deloitte.com).

Eli Lilly's approach to this phase illustrates the capital intensity that full-scale AI infrastructure can require at the largest companies. In October 2025, Lilly announced it was deploying what it described as the world's largest AI factory wholly owned and operated by a pharmaceutical company, built with over **1,000** NVIDIA Blackwell Ultra GPUs delivering more than **9,000** petaflops of AI performance, part of a broader **$50 billion** commitment to expanding its U.S. manufacturing and R&D footprint (blogs.nvidia.com). Select models trained on this infrastructure are being made available externally through Lilly's TuneLab platform, built on **$1 billion** worth of Lilly's proprietary data, giving smaller biotech companies a path to advanced AI capability without building comparable infrastructure themselves (blogs.nvidia.com). The scale of that leap is easiest to grasp by comparison: Lilly's own framing notes that in 1992 its Cray supercomputer represented the pinnacle of scientific computing, while a single Blackwell Ultra GPU in the new AI factory now contains roughly the computing power of about **7 million** of those Cray systems, delivering over **9 quintillion** math operations per second across the full 1,016-GPU cluster (blogs.nvidia.com). Lilly's own chief AI officer, Thomas Fuchs, has framed the shift as strategic rather than purely technical, saying the company is "shifting from using AI as a tool to embracing it as a scientific collaborator" (prnewswire.com), consistent with treating AI infrastructure as a multi-year capability build (prnewswire.com), one built, as Lilly executives note, on the "decades of data" the 150-year-old company already holds (prnewswire.com), running the new supercomputer on the world's first NVIDIA DGX SuperPOD built with DGX B300 systems (prnewswire.com), part of an AI lifecycle platform spanning data ingestion and training through fine-tuning and high-volume inference (prnewswire.com).

Implementation Considerations and Process Changes

Executing the four-phase timeline above requires attention to several cross-cutting process changes that determine whether an organization actually reaches enterprise scale on schedule. The first is **data foundations**. Multiple industry analyses converge on the finding that neglecting data quality and governance early is the leading cause of AI initiative failure in biotech, since AI models trained on fragmented, poorly annotated, or non-FAIR (Findable, Accessible, Interoperable, Reusable) data cannot be validated to the standard the FDA's risk-based credibility framework expects for regulatory-facing use cases (fda.gov).

The second is **build-versus-buy discipline**. BCG's guidance is direct on this point: choosing whom to work with, where, and in what capacity requires determining up front which capabilities should be built in-house and which should be sourced externally, since AI-based drug discovery startups are already partnering extensively with established biopharma companies to accelerate R&D rather than compete with it (bcg.com). Bessemer Venture Partners has observed that the vendor ecosystem serving pharma, contract research organizations (CROs) billing by headcount, platform vendors licensing per seat, has structural incentives misaligned with AI-driven automation, since automating a workflow can threaten the very unit economics that fund the vendor relationship (bvp.com). This makes vendor selection during roadmap execution as much a strategic decision as a procurement one (intuitionlabs.ai); a biotech that outsources its AI roadmap entirely to a vendor with disincentives to automate is unlikely to hit an aggressive scaling timeline (fiercepharma.com).

The third is **training and change management at scale**, which becomes the binding constraint once a pilot succeeds and needs to reach thousands of end users. Merck's internal generative AI platform, known internally as GPTeal, illustrates the scale of this challenge: more than **50,000** Merck employees, out of a total workforce of roughly 75,000, were using the platform regularly as of early 2025, a rollout supported by self-service e-learning modules, monthly webcast seminars, and developer bootcamps ranging from half a day to ten days in length (intuitionlabs.ai). In a collaboration with McKinsey's QuantumBlack unit, Merck used generative AI to cut the time required to produce draft clinical study reports from the usual two to three weeks down to **three to four days**, while human-reviewed drafting hours per report fell from about 180 to 80 and error rates in those drafts were cut roughly in half.

The fourth is **regulatory sequencing**. Because the EU AI Act imposes statutory dates and documentation requirements independent of a company's internal roadmap timing, and the FDA's draft AI guidance sets out nonbinding recommendations that sponsors should anticipate, biotech companies need to align their scaling phase with these external milestones rather than treat compliance as a parallel workstream to be addressed later. The EU AI Act's Annex III high-risk system obligations now apply from 2 December 2027 (and Annex I product-embedded systems from 2 August 2028) under the Digital Omnibus on AI, Regulation (EU) 2026/1744, so companies are planning around settled law rather than a proposal (eur-lex.europa.eu). Article 11 of the Act requires providers of high-risk AI systems to draft and maintain comprehensive technical documentation, covering system design, training data governance, and accuracy and robustness testing, before placing a system on the market, work that mature governance practices need to have produced well before any regulatory deadline arrives (a-lign.com). The EMA's own reflection paper on AI in the medicinal product

lifecycle, opened for public consultation covering the period from July to December 2023, likewise recommends that any AI or machine learning system expected to affect a medicine's benefit-risk balance be brought to regulators early, through mechanisms such as qualification of innovative development methods or formal scientific advice, rather than only at the point of marketing authorization (ema.europa.eu) (reuters.com).

The fifth is **choosing and evolving an organizational model**, since the centralized, decentralized, and hybrid archetypes each carry different speed and control tradeoffs. In a centralized model, decisions are driven top down, with the C-suite pushing teams to integrate analytics and AI throughout drug development; in a decentralized model, domain-specific insights and service delivery speed take priority over central standardization (zs.com). Organizations pursuing any of the three models are also increasingly building explicit AI ethics commitments into their governance layer; Novartis, for example, has outlined eight principles for the ethical and responsible use of AI, including accountability, bias mitigation, privacy, transparency, and environmental sustainability, that inform how its AI Center of Excellence evaluates new use cases (zs.com). Industry-wide talent pooling has also emerged as a Phase 3 accelerant: the Machine Learning for Pharmaceutical Discovery and Synthesis Consortium (MLPDS), facilitated by MIT, brings together **12** leading biopharma companies around a shared interest in AI-driven discovery and synthesis tools, letting member companies shorten their own capability build-out by pooling research investment rather than duplicating it independently (zs.com). ZS's research also identifies talent retention and community-building as a distinct dimension of Phase 3 execution: leading organizations pursue purposeful data strategies through acquisitions and partnerships, establish themselves as destinations for data science talent through hackathons, symposiums, and conferences, and tie performance incentives to innovation while offering job rotations that keep top performers engaged across functions (zs.com).

Table 1 below compares the three organizational archetypes that recur across the sources reviewed in this report, since the choice of model materially affects how quickly a biotech can move from Phase 3 pilot selection into Phase 4 scaling.

Organizational Model	Decision Authority	Typical Strength	Typical Weakness
Centralized	C-suite and central AI/data science function drive priorities top down	Consistency, easier regulatory documentation, economies of scale on infrastructure	Slower to reflect domain-specific nuance in R&D sub-functions
Decentralized	Domain teams (discovery, clinical, manufacturing) own their own AI priorities	Fast, close to the problem, high domain relevance	Duplicated tooling, inconsistent governance, harder to scale enterprise-wide
Hybrid (Hub-and-Spoke)	Central hub sets data strategy and shared capability; domain "spokes" apply it locally	Balances consistency with domain relevance; increasingly the prevalent model among large biopharma	Requires strong coordination discipline between hub and spokes to avoid drift

As *Table 1* shows, the hybrid model's advantage is not that it eliminates tradeoffs but that it lets a company sequence them: the central hub can absorb the up-front cost of building shared infrastructure and governance during Phase 3, while domain spokes iterate quickly on their own use cases without waiting for every capability to be centrally built first, which is consistent with why ZS's benchmarking study found this model becoming the prevalent choice among the large and midsize biopharma companies it studied.

Data Analysis and Evidence

The quantitative picture of biotech AI adoption and timelines is more mixed than headline enthusiasm statistics suggest, and the discrepancies are informative in their own right. Reuters reporting on industry-wide AI partnerships from 2025 through mid-2026 found that industry forecasts suggest machine learning applied to target discovery, molecule design, and clinical trial planning could halve early-stage development timelines and costs within three to five years (reuters.com), a deal cadence including Bristol Myers Squibb's May 2026 deployment of Anthropic's Claude to more than **30,000** employees (reuters.com) and Novo Nordisk's OpenAI partnership spanning discovery through commercial operations (reuters.com). On adoption intent, the numbers are consistently high: Capgemini's Research Institute found **82%** of biopharma executives believe AI will fundamentally transform R&D and **60%** agree that companies failing to scale AI will fall behind competitively, while **79%** are actively formalizing AI integration strategies (capgemini.com). A separate survey of the industry's largest firms, health-tech venture capital (VC) firm Define Ventures' report drawing on conversations with executives from 16 of the top 20 ranked pharmas, found **70%** of pharma leaders call AI an "immediate priority," rising to **85%** among the top-ranking Big Pharmas specifically, with more than **80%** of respondents increasing AI budgets and only **15%** holding budgets flat (fiercepharma.com).

On realized outcomes, the numbers are far lower. Deloitte's 2026 survey found just **22%** of life sciences leaders had successfully scaled AI and only **9%** reported significant returns, as noted above, and ZS's benchmarking of 15 large and midsize biopharma companies found a nearly identical **20%** successful-scaling rate, as detailed in Phase 4 above. This adoption-versus-scaling gap, roughly 60 to 85 percentage points of intent unmatched by realized value, is the single most important data point for anyone building a biotech AI timeline: it indicates the constraint is not enthusiasm, technology availability, or even budget, but the organizational and process discipline required to move from pilot to production described in Phase 4 above.

Table 2 below summarizes the phased timeline described in this report, mapping each phase to its typical duration, primary deliverable, and the most commonly cited failure mode at that stage, drawn from the sources referenced throughout this analysis.

Phase	Typical Duration	Primary Deliverable	Most Common Failure Mode
1. Diagnostic and Readiness Assessment	60 to 90 days	AI system inventory, stakeholder map, data and governance gap analysis	Incomplete inventory misses shadow AI use, understating true governance exposure
2. Strategy Formulation and Board Approval	3 to 6 months	Board-approved strategy with financial model, use-case prioritization matrix	Strategy lacks financial rigor comparable to other capital requests, triggering rework
3. Governance, Team Formation, and Pilot Selection	6 to 12 months	AI governance body, hub-and-spoke or CoE org design, 5 to 10 funded pilots	Too many parallel pilots dilute funding and attention (bcg.com)
4. Scaling and Enterprise Integration	Year 2 to Year 3	Enterprise-wide deployment, measurable ROI, integrated data and MLOps infrastructure	Pilot-to-production transition stalls; only 20 to 22% of firms report successful scaling

The pattern in *Table 2* is that duration alone does not predict success; the failure modes cluster around discipline (scope, financial rigor, data quality) rather than around calendar time. A biotech that spends longer on Phase 1 and Phase 2 but arrives at Phase 3 with a genuinely prioritized, well-governed portfolio of pilots appears, across the sources reviewed here, to reach durable Phase 4 scaling faster than one that rushes into dozens of parallel pilots without that discipline.

A systematic peer-reviewed review published in *Pharmaceuticals* and indexed on PubMed Central adds a scientific-literature perspective to these industry-survey figures. Analyzing 173 studies published between 2015 and 2025, the review found that machine learning (ML) methods accounted for **40.9%** of AI techniques used in drug discovery and development, with molecular modeling and simulation at **20.7%** and deep learning (DL) at **10.3%**, and that oncology was by far the most-studied therapeutic area at **72.8%** of studies, followed by dermatology at **5.8%** and neurology at **5.2%** (pmc.ncbi.nlm.nih.gov). On clinical translation, the same review found **39.3%** of studies concentrated at the preclinical stage and **23.1%** in Clinical Phase I, with only **45%** of studies reporting measurable clinical outcomes and **97%** reporting formal industry partnerships, evidence that AI's documented impact remains concentrated in early discovery rather than late-stage clinical proof, a pattern any realistic biotech timeline should anticipate (pmc.ncbi.nlm.nih.gov). Formal industry partnerships were in fact the norm rather than the exception across the full sample: **168** of the 173 studies reviewed disclosed a formal industry partnership, reinforcing the broader pattern in this report that even AI-native biotechs rarely execute their timelines in isolation from larger pharmaceutical partners (pmc.ncbi.nlm.nih.gov).

Market-level figures reinforce the scale of what is at stake in getting the timeline right. Bessemer Venture Partners estimates the addressable outsourced pharma and life sciences services market at over **$150 billion**, with generative and agentic AI capable of unlocking over **$100 billion** in net new annual value across the end-to-end pharma lifecycle, while noting that only about **5%** of pharma companies have captured measurable AI value so far due to legacy data silos, fragmented ownership, and vendor models built on billable headcount (bvp.com). Separately, BCC Research projects annual investment in AI-driven drug discovery to grow from **$3.8 billion** in 2025 to **$15.2 billion** by 2030, a roughly fourfold increase over five years, cited by GSK's own leadership as validation of the direction the industry is heading (govinfosecurity.com). Bessemer breaks the value unlock into three reinforcing levers that map onto the phased timeline in this report: cost reduction from fewer failed programs and leaner spend, faster time-to-market from compressed development cycles, and higher launch-year revenue from better commercial targeting, arguing that AI expands the total economic pie rather than simply reallocating the existing $150 billion in outsourced spend (bvp.com).

Two further data points round out the picture on where in the discovery pipeline AI activity actually concentrates and how firms are financing it. Capgemini's Research Institute found that **74%** of executives specifically believe generative AI holds significant potential in drug discovery, a figure consistent with the peer-reviewed finding that preclinical and discovery-stage work accounts for the largest single share of published AI applications (capgemini.com). On the regulatory side, the FDA has said it will continue building out formal external engagement pathways specifically for companies developing AI in drug development, on top of the informal channels, workshops, and comment periods that already shaped its January 2025 draft guidance (fda.gov), a signal that companies building a multi-year roadmap should plan for iterative regulatory dialogue rather than a single guidance document that will remain static throughout their scaling phase.

Case Studies and Real-World Examples

Amgen: A Chief Technology Officer Mandate Built Around "The Hinge Moment"

In 2023, Amgen's then-head of R&D, Dave Reese, made the case to his chief executive officer (CEO) that the company needed a newly created chief technology officer (CTO) role to lead an AI-driven transformation, a pitch he later described internally as recognizing "the hinge moment" for AI in biotech (endpoints.news). Two years on, Amgen's own scientific leadership describes the resulting transformation as continuous rather than complete: there was no single breakthrough moment, but instead an ongoing reorganization of large parts of the biopharmaceutical company around AI models embedded increasingly deeply into the discovery process. Amgen's case is instructive precisely because it does not claim to have "finished" its AI roadmap; it illustrates that even a company that moved decisively on organizational structure in year one is still, years later, treating the scaling phase as ongoing work rather than a completed milestone.

GSK: Six Years of Infrastructure Before the Headlines

GSK offers the clearest documented example of the multi-year timeline this report describes. When a journalist visited GSK's London research base earlier in 2026, the correspondent found no lab benches, no equipment, and no white coats; the researchers observed were instead building software trained to read genomes and generate disease hypotheses entirely inside a computer, a visible sign of a transformation that had been underway for years by the time it became publicly visible (govinfosecurity.com). According to reporting on the company's transformation, GSK has spent the past **six years** quietly rebuilding its research and development operating model around AI, from foundational data infrastructure through to AI agents now capable of writing scientific hypotheses (govinfosecurity.com). One tangible output of that investment is Undermind, an AI system that exhaustively follows citation trails to synthesize scientific literature; more than **1,000** GSK scientists now use it daily. In an internal survey of over **130** of its scientists, the tool achieved a Net Promoter Score (NPS) of **63**, described as well above published benchmarks for enterprise software and placing it in the 95th percentile (govinfosecurity.com). GSK has also built a combined imaging and omics analysis platform whose high-content imaging component is at least **50 times** more efficient than standard omics profiling, and which in one colorectal cancer study predicted approximately **30%** of the transcriptome from imaging data alone while recapturing **98%** of key findings using only **30%** of the usual laboratory experiments (govinfosecurity.com). In a single day in January 2026, GSK signed two AI-focused external deals, including a multi-year data partnership with genomics firm Helix and a five-year, **$50 million** agreement with AI-native biotech Noetik for virtual cell models in oncology, illustrating how a mature AI program continues to combine internal build with external partnership even after years of internal investment (govinfosecurity.com).

Moderna: Scaling Without Scaling Headcount

Moderna's AI strategy is organized around a specific structural constraint rather than a general enthusiasm for the technology. The company reported approximately **$1.9 billion** in revenue for 2025, down sharply from prior years as COVID-19 vaccination rates normalized, while research and development spending totaled roughly **$3.1 billion**, and it employs around **4,700** people, a workforce it has explicitly chosen not to scale in proportion to its product ambitions (emerj.com). Moderna executives have argued that bringing the roughly **15** new products the company hopes to launch over five years to market under a traditional operating model would require a headcount in the hundreds of thousands, a gap AI is meant to close by multiplying the output of its existing workforce rather than growing it (emerj.com). After comparative testing against Microsoft Copilot, Moderna selected ChatGPT Enterprise as its company-wide platform, and within about two months of launch employees had built more than **750** custom GPTs tailored to their own teams' documents and terminology without engineering support (emerj.com). Moderna's legal department has reached **100%** adoption of ChatGPT Enterprise, the highest of any function company-wide, with platform access expanding from an initial rollout of roughly 3,000 employees to deployment across legal, research, manufacturing, and commercial functions enterprise-wide (emerj.com). Moderna's own framing of the strategy is that AI lets the company operate like a much larger organization without a proportional headcount increase, though the company itself treats that framing as a directional benefit rather than an audited financial result, since sustained usage metrics such as full legal-department adoption are a real signal of embedding even before formal return-on-investment figures are disclosed (emerj.com).

Merck: GPTeal and the Economics of Enterprise-Wide Rollout

Merck & Co. (known as MSD outside North America) built a proprietary, gated large language model (LLM) interface called GPTeal that lets employees interact securely with models including OpenAI's ChatGPT, Meta's Llama, and Anthropic's Claude from inside a controlled corporate environment. As of early 2025, more than **50,000** Merck employees, out of a total workforce of roughly 75,000, were actively using the platform regularly (intuitionlabs.ai). In its most measurable application, a partnership with McKinsey's QuantumBlack unit used generative AI to cut clinical study report drafting time from two to three weeks down to three to four days, cutting human-reviewed drafting hours per report from roughly 180 to 80 and reducing error rates in those drafts by approximately half. By June 2025, Merck confirmed that the first live AI-assisted clinical study reports had been submitted and that the platform was scaling across its late-phase pipeline, illustrating the transition from Phase 3 pilot to Phase 4 enterprise scaling described earlier in this report. Merck's rollout also illustrates that there is no single correct posture toward employee AI adoption: Merck's own comparative positioning places it between Johnson & Johnson's more top-down mandated training model and Eli Lilly's comparatively open, encouragement-led approach that pairs free access with a formal internal AI certification program, evidence that the training and change-management workstream in Phase 3 and Phase 4 can be sequenced several different ways depending on a company's existing culture (intuitionlabs.ai).

Insilico Medicine and Exscientia: AI-Native Biotechs Compressing Preclinical Timelines

Where the large-pharma cases above show incumbents retrofitting AI into existing organizational structures, AI-native biotechs illustrate what a timeline built around AI from inception can achieve. Insilico Medicine announced in 2021 that it had identified a novel target for idiopathic pulmonary fibrosis and advanced a drug candidate, later designated INS018_055, from target identification to preclinical candidate in under **18 months**, a process that traditionally takes four to six years, at a disclosed cost of only **$150,000** excluding wet lab validation (pmc.ncbi.nlm.nih.gov). That candidate has since entered Phase II clinical trials, underscoring how AI can compress not just discovery timelines but the bridge between computational work and clinical translation (pmc.ncbi.nlm.nih.gov). Exscientia (since November 2024 a wholly owned subsidiary of Recursion), in partnership with Sumitomo Dainippon Pharma, developed a novel small-molecule drug candidate, DSP-1181, for obsessive-compulsive disorder in less than **12 months**, making it the first AI-designed molecule to enter human clinical trials (pmc.ncbi.nlm.nih.gov). Recursion Pharmaceuticals represents a third distinct model, using automated high-throughput imaging combined with deep learning to identify phenotypic changes in cells, enabling rapid repurposing of existing molecules alongside discovery of novel therapeutics (pmc.ncbi.nlm.nih.gov).

AstraZeneca: Embedding AI Across the Full R&D Value Chain Rather Than a Single Function

AstraZeneca's case illustrates a fifth pattern: rather than centering its timeline on one signature platform or one department, the company describes applying AI "throughout the discovery and development process, from target identification to clinical trials," treating AI as infrastructure embedded across the full R&D value chain rather than a bolt-on capability in a single function (astrazeneca.com). On the clinical trials side specifically, the company points to the rapid adoption of high-quality electronic health records (EHRs) as a data source with substantial potential to improve trial implementation, and describes federated EHR technology as unlocking new opportunities to refine or replace clinical trial processes including patient identification, selection, and data capture (astrazeneca.com). AstraZeneca also illustrates the governance dimension of a mature AI timeline: during 2020 the company engaged a diverse range of internal and external experts specifically to develop principles for the ethical and responsible use of AI, aligned to its existing code of ethics, well before generative AI's mainstream adoption made such governance a competitive necessity for the rest of the industry (astrazeneca.com). That five- to six-year head start on governance, well before the FDA's 2025 draft guidance or the EU AI Act's high-risk obligations existed, is consistent with the broader finding in this report that the companies furthest along in AI scaling today are the ones that treated governance as a Phase 2 deliverable rather than a Phase 4 afterthought. The company frames the stakes of that early investment directly: it notes that more data has been generated in the past two years alone than in the entire previous history of the human race, a scale argument it uses internally to justify building AI and data science capability as core R&D infrastructure rather than a specialist side function (astrazeneca.com).

Implications and Future Directions

The clearest implication of the evidence assembled in this report is that biotech leaders should stop asking "how long will our AI strategy take" as a single question and start planning against the four distinct phases identified here, each with its own duration, deliverable, and characteristic failure mode. Companies that collapse the diagnostic and strategy-formulation phases to save time often produce documents that boards send back for rework; in the author's judgment, companies that skip disciplined pilot prioritization in favor of proliferating proofs of concept often fail to reach the scaling phase at all. The roughly 60-percentage-point gap between AI adoption intent and realized scaled value, visible consistently across Deloitte, ZS, and other benchmarking sources cited throughout this report, is the strongest available evidence that timeline discipline, not technology access, is now the binding constraint on biotech AI value creation.

Regulatory timing will increasingly compress the available planning window rather than expand it. With the FDA's AI credibility framework already informing submission expectations and the EU AI Act's Annex III high-risk obligations applying from 2 December 2027 under the Digital Omnibus on AI, as detailed above, biotech companies beginning a strategy process in 2026 or later effectively need to build regulatory documentation workstreams into Phase 2 rather than treating them as a Phase 4 afterthought. Gartner's broader enterprise research suggests the workforce implications of this shift are already visible outside biotech specifically: by 2030, the firm's July 2025 survey of CIOs found an expectation that **75%** of information technology work will be done by humans augmented by AI and **25%** by AI alone, with **0%** expected to be done by unaugmented humans (gartner.com). Whether biotech R&D and commercial functions follow that same trajectory on a similar timeline remains an open question, but the direction of travel among the industry's most digitally mature companies, GSK, Merck, Sanofi, Moderna, and Lilly among them, points the same way. Gartner itself frames this as workforce transformation rather than pure headcount reduction, predicting AI will create more jobs than it destroys by 2028 (gartner.com). Regulators are also moving to coordinate their own AI oversight rather than leaving it fragmented across product centers: the FDA published a cross-center document in March 2024, revised in February 2025, describing how its drug, biologics, device, and combination-product centers plan to align their AI oversight approaches so that, to the extent feasible, common regulatory approaches can apply across different medical product types (fda.gov), a signal that a biotech's regulatory affairs function should expect increasingly harmonized, rather than siloed, AI guidance as its roadmap matures.

A second implication concerns organizational design. The hub-and-spoke or Center of Excellence model that ZS's benchmarking study and IntuitionLabs' organizational design analysis both describe as the emerging norm is likely to keep evolving rather than settle into a fixed template, particularly as agentic AI, systems capable of autonomous multi-step action rather than single-response generation, moves from the 30% of firms currently citing it as an area of interest, as noted above, toward broader production deployment. Companies building their AI Center of Excellence today should design the governance layer to be extensible to agentic and multi-agent systems rather than assuming today's single-model chatbot governance framework will remain adequate through the full scaling phase.

Finally, the role of external advisors in this timeline deserves honest treatment. Given the documented gap between strategy formulation and successful scaling, life sciences consultancies and specialized AI advisory firms are increasingly positioned not as AI technology vendors competing with platform providers, but as

advisors helping biotech and pharma clients navigate exactly the sequencing, governance, and regulatory-alignment questions this report has described, particularly where that advisory work intersects with existing enterprise content and quality systems such as Veeva Vault, where the aim is to connect AI initiatives to systems clients already operate rather than displace them.

Frequently Asked Questions (FAQs)

How long does it take to build a biotech AI strategy? As a planning heuristic synthesized by the author rather than a measured industry benchmark, a diagnostic and readiness assessment typically takes 60 to 90 days, strategy formulation and board approval an additional three to six months, governance and initial pilot selection another six to twelve months, and full enterprise scaling a further one to two years, for a total of roughly two to three years from a standing start to a mature, enterprise-scaled AI capability, noting that the multi-year build-outs described at GSK and Sanofi in the Case Studies section above illustrate the scale of effort but do not themselves validate this total (govinfosecurity.com).

What is the difference between a biotech AI strategy and a biotech AI roadmap? A strategy identifies where AI will create durable competitive advantage and how the organization will be structured to capture it; a roadmap is the sequenced, resourced implementation plan, with named owners, milestones, and budget, that turns that strategy into deployed capability. Confusing the two is a common cause in practice of AI initiatives that receive board approval but never reach production.

How much does building a biotech AI department or Center of Excellence cost? Public disclosures vary enormously by company scale, from Sanofi's Toronto hub growing from 2 to 150 people over two years (jobs.sanofi.com) to Eli Lilly's infrastructure investment tied to a broader **$50 billion** manufacturing and R&D commitment (blogs.nvidia.com). Gartner recommends organizations plan for roughly **10** hidden costs beyond licensing for every AI tool purchased, including training and change management (gartner.com).

Should a biotech build AI capability in-house or hire consultants? Most evidence points toward a hybrid answer, as detailed in the Implementation Considerations section above: BCG recommends determining explicitly which capabilities should be built internally versus sourced externally rather than defaulting to either extreme, and Bessemer Venture Partners separately cautions that vendors whose economics depend on billable headcount or per-seat licensing have structural disincentives to help a client automate away that spend (bvp.com).

What is an AI Center of Excellence in biotech, and does every company need one? A CoE is a centralized entity that orchestrates AI capability, standardizes best practices, and provides oversight for data quality and regulatory compliance across an organization (intuitionlabs.ai); as detailed above, hybrid hub-and-spoke models, where a central CoE builds shared capability and domain teams apply it, were the most common structure among the biopharma companies ZS studied in its benchmarking research.

How long does it take to see return on investment from biotech AI? Evidence suggests ROI is uneven and often slower than expected: only 22% of life sciences leaders in Deloitte's 2026 survey reported successfully scaling AI and just 9% reported significant returns, as cited above, even though 80% of top pharma companies report increasing AI budgets, as noted earlier in this report.

Does the FDA's AI guidance change the biotech AI timeline? Yes, for any AI model whose output will support a regulatory submission. The FDA's January 2025 draft guidance recommends a risk-based credibility assessment tied to a model's specific context of use, meaning validation and documentation work for regulatory-facing AI needs to start well before a submission date rather than being retrofitted afterward (fda.gov). The agency has signaled this is an evolving area rather than a one-time document: CDER has committed to continuing to develop and adopt a risk-based regulatory framework for AI in drug development as the technology and its own experience with submissions mature (fda.gov).

How does the EU AI Act affect a biotech AI roadmap? Any biotech deploying a high-risk AI system, as defined under Annex III, needs technical documentation, training-data governance records, and accuracy and robustness testing artifacts in place before the applicable high-risk application date, which under the Digital Omnibus on AI (Regulation (EU) 2026/1744) is 2 December 2027 for Annex III systems and 2 August 2028 for Annex I product-embedded systems (eur-lex.europa.eu). Not every biotech AI system is high-risk, so classifying each system by its intended use is an early roadmap task. Companies may also consider an AI management system such as ISO 42001 as a governance foundation.

Conclusion

The question "how long does it take to build a biotech AI strategy timeline" resolves, once strategy and roadmap are properly separated, into a phased answer supported by consistent evidence across regulatory bodies, management consultancies, and named corporate case studies: 60 to 90 days for an honest diagnostic, three to six months for a board-approvable strategy, six to twelve months to establish governance and launch a disciplined set of pilots, and one to two additional years to reach genuine enterprise scale. Companies that have executed this arc successfully, GSK over six years, Sanofi's Toronto hub over two years, Merck's GPTeal rollout to more than 50,000 employees, share a common pattern of sustained, disciplined investment rather than a single rapid transformation event.

The data equally makes clear that most organizations are still earlier in this journey than their public statements suggest. With adoption intent running as high as 82% to 85% among top pharmaceutical executives but successful scaling reported by only 20% to 22% of organizations, the gap between ambition and execution remains the defining feature of the current market. Regulatory deadlines, the FDA's risk-based credibility framework (still nonbinding draft guidance) and the EU AI Act's Annex III obligations applying from 2 December 2027, will not wait for that gap to close on its own. Biotech and pharmaceutical leaders building an AI roadmap today should treat the phased timeline described in this report not as an aspirational best case, but as the realistic minimum against which to plan governance, budget, and regulatory readiness, informed by advisors who understand both the technology and the regulated environment in which it must operate.

Source links

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

1. https://www.bcg.com/publications/2023/biopharma-path-to-value-with-generative-ai
2. https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html

3. https://www.capgemini.com/wp-content/uploads/2026/01/Final-Web-Version-Report-Gen-AI-in-Lifesciences.pdf
4. https://www.bvp.com/atlas/roadmap-reinventing-life-sciences-with-ai
5. https://jobs.sanofi.com/en/inside-sanofi-ai-engine
6. https://endpoints.news/in-the-lab-how-amgen-sees-ai-changing-pharma-rd
7. https://www.govinfosecurity.com/blogs/gsk-ai-driven-science-factory-p-4121
8. https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development
9. https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
10. https://pmc.ncbi.nlm.nih.gov/articles/PMC12298131
11. https://blogs.nvidia.com/blog/lilly-ai-factory-nvidia-blackwell-dgx-superpod
12. https://www.gartner.com/en/newsroom/press-releases/2025-11-10-gartner-survey-finds-artificial-intelligence-will-touch-all-information-technology-work-by-2030
13. https://www.fiercepharma.com/marketing/ai-tech-immediate-priority-most-big-pharmas-many-plan-open-their-pockets-further-ai
14. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological
15. https://www.ema.europa.eu/en/news/reflection-paper-use-artificial-intelligence-lifecycle-medicines
16. https://www.reuters.com/business/healthcare-pharmaceuticals/pharma-sector-doubles-down-ai-amid-hopes-slashing-costs-timelines-2026-05-20
17. https://www.zs.com/insights/how-to-use-ai-in-pharma-r-and-d-artificial-intelligence-biopharma
18. https://www.astrazeneca.com/r-d/data-science-and-ai.html
19. https://www.prnewswire.com/news-releases/lilly-partners-with-nvidia-to-build-the-industrys-most-powerful-ai-supercomputer-supercharging-medicine-discovery-and-delivery-for-patients-302597285.html
20. https://www.a-lign.com/articles/eu-ai-act-enforcement-delay
21. https://emerj.com/ai-at-moderna

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.