Biotech AI Budget Benchmarks 2026: Mid-Market Spending Data

6/24/2026 • 45 min read

biotech ai budget | ai implementation cost | life sciences ai spending | biotech ai adoption | ai center of excellence pharma | drug discovery roi | biotech it budget | mid-market biotech ai

Executive Summary

Mid-market biotech and pharmaceutical companies enter the second half of 2026 with no single accepted benchmark for artificial intelligence (AI) spending, but a converging set of data points now makes reasonable planning possible. Life sciences technology leaders report that **97% of organizations allocate less than 1% of annual revenue to digital technology** overall ([1] assets.kpmg.com), even as **87%** of the same organizations already report AI agents embedded in workflows, products, or services ([2] assets.kpmg.com). Within research and development (R&D) budgets specifically, Benchling's November 2025 survey of roughly 100 biotech and pharma organizations found that **55% of respondents devote 11% or more of their R&D tech budget to AI**, split across bands of 11 to 20% (28%), 21 to 35% (18%), and above 35% (6%) ([3] www.benchling.com). At the enterprise IT level, healthcare and life sciences organizations broadly spend **6 to 9% of revenue on IT**, above the 5.7% cross-industry average ([4] itbudgetcalculator.com), while mid-market companies overall are budgeting **2 to 5% of revenue for AI specifically**, concentrated at 2.5 to 3.5% ([5] aisavvy.io).

For a concrete anchor at the deal-size level that mid-cap biotech leaders actually operate under, industry advisors put a realistic annual AI Center of Excellence (CoE) budget for a **$500 million to $5 billion revenue mid-cap pharma company at approximately $2 million**, split roughly 60% staffing, 20% technology and tooling, 12.5% external partnerships, 5% training, and 2.5% experimentation reserve ([6] sakaradigital.com). By contrast, large-cap pharma enterprise AI investment runs **$50 million to $200 million annually**, roughly ten times the mid-cap envelope, as detailed further below. These figures sit inside a much larger capital backdrop: global AI spending across all industries is forecast to reach **$2.52 trillion in 2026**, a 44% year-over-year increase ([7] www.gartner.com), and worldwide IT spending is projected at **$6.15 trillion**, up 10.8% ([8] www.gartner.com).

Timelines for going from pilot to measurable production value remain long relative to enterprise software norms. Deloitte's 2026 Life Sciences Outlook Survey found that only **22% of life sciences leaders say they have successfully scaled AI**, and just **9% report achieving significant returns** on those efforts ([9] www.deloitte.com). ZS's 2026 survey of 115 pharma and biotech technology executives similarly found that only **17% can prove measurable AI value in drug discovery today**, though 42% expect it within a year, and only about **40% of AI pilots that get funded eventually reach scaled deployment**, as detailed further below. Infrastructure costs are falling on the compute side even as headline AI budgets grow: Amazon Web Services (AWS) cut P5 H100 graphics processing unit (GPU) instance list pricing by 44% in mid-2025, bringing on-demand cost to roughly **$6.88 per GPU-hour**, with alternative cloud providers now offering the same hardware from under $2 per GPU-hour (www.spheron.network).

Case evidence supports the ROI case even where scaling remains incomplete. Pfizer's collaboration with AWS saved scientists **16,000 hours of search time annually** and cut infrastructure costs by **55%** ([10] aws.amazon.com). Roche's newly expanded NVIDIA-powered AI factory now exceeds **3,500 Blackwell GPUs**, the largest disclosed pharmaceutical GPU footprint ([11] www.roche.com). Recursion Pharmaceuticals, a clinical-stage "TechBio" company, ended 2025 with **$399.2 million in cash operating expense** and reported delivering development candidates by synthesizing roughly **330 compounds per program over about 17 months**, versus an industry average of over 2,500 compounds and 42 months (ir.recursion.com). This report walks through the budget benchmarks, implementation timelines, infrastructure cost drivers, and governance frameworks that mid-market biotech leaders need to plan a defensible AI enablement strategy, without overstating what the current data can support.

Introduction and Background

Biotech and pharmaceutical executives asking "how much should we spend on AI" confront an unusually noisy data environment. Large-cap pharma companies publish investor-facing AI announcements measured in hundreds of millions or billions of dollars, while independent surveys of the broader life sciences sector consistently show the median organization spending under 1% of revenue on digital technology as a whole ([13] assets.kpmg.com). For a mid-market biotech company, typically defined here as one with revenues between roughly $500 million and $5 billion or a smaller, venture-backed clinical-stage company with a correspondingly modest operating budget, neither the large-cap headline numbers nor the cross-industry averages provide a usable planning figure. This report assembles the available benchmark data as of **July 2026** into a single reference: what percentage of revenue or R&D budget peer organizations allocate to AI, how long implementation actually takes from pilot to production, what infrastructure and talent costs look like at current market rates, and what a defensible mid-market AI enablement strategy contains.

The urgency behind this question is not abstract. KPMG's 2026 Global Tech Report, based on a survey of **124 Life Sciences technology function leaders** at large and mid-sized pharmaceutical, biotechnology, and medical device companies, found that **48% of organizations have integrated AI strategically into core business functions**, while **44% report limited maturity in funding, supporting, or scaling AI and automation** ([14] assets.kpmg.com) ([15] assets.kpmg.com). That gap between adoption enthusiasm and funding discipline is precisely where mid-market organizations get into trouble: they either underinvest and fall behind on foundational data infrastructure, or overcommit to headline-grabbing pilots that never reach production and drain scarce capital. The same KPMG survey found that **52% of life sciences organizations identify as "fast followers,"** adopting new technologies only after validation by early movers, while a smaller but meaningful **35% self-identify as innovators** willing to lead and experiment early ([16] assets.kpmg.com)([17] assets.kpmg.com), a split mid-market leaders should weigh honestly when deciding how aggressively to fund AI ahead of proven peer results.

Drug development economics compound the stakes. Historically, bringing a single new medicine to market has required **typically 10 years at a cost of US$2.5 billion** ([18] docs.oracle.com), with roughly **90% of drug candidates that enter Phase 1 trials never reaching patients** ([19] www.drugpatentwatch.com). AI's promise is to compress timelines and reduce attrition-driven waste inside that math, but only if implementation dollars translate into deployed, validated systems rather than stranded pilots. This report is organized around four questions mid-market biotech leaders consistently ask: What do peers actually spend, as a percentage of revenue and R&D budget? How long does implementation realistically take? What does the infrastructure and talent cost stack look like at 2026 market rates? And what does a credible, board-defensible AI enablement strategy include, given that governance and regulatory expectations (from the U.S. Food and Drug Administration, or FDA, and the European Medicines Agency, or EMA) are now explicit rather than implicit.

As an adjacent advisory perspective grounded in life sciences technology and Veeva ecosystem work, IntuitionLabs.ai's own analysis of pharma and biotech AI adoption trends frames the sector's current stage as one where **AI adoption is already widespread at the pilot level but few organizations have scaled to enterprise-wide programs** ([20] intuitionlabs.ai), a framing borne out repeatedly in the benchmark data assembled below.

Budget Benchmarks: What Biotech Companies Actually Spend on AI

Three distinct budget lenses matter for mid-market planning, and they should not be conflated: enterprise IT spend as a percentage of revenue, the share of R&D technology budget specifically allocated to AI, and absolute dollar envelopes for defined organizational functions like an AI Center of Excellence.

At the broadest level, healthcare and life sciences organizations spend **6 to 9% of revenue on total IT**, compared with a **5.7% average across all industries** in 2026, with financial services and technology running slightly higher at 7 to 10% and manufacturing or logistics running lower at 2 to 5% ([21] itbudgetcalculator.com). Company size shifts this ratio meaningfully: smaller organizations spend a higher percentage of revenue on IT (roughly 6.9% for startups) while large enterprises spend a lower percentage (around 3.7%) due to economies of scale, with mid-market companies (251 to 1,000 employees) typically falling between **6 and 8% of revenue** ([22] itbudgetcalculator.com). Per-employee IT spend averages roughly **$9,000 to $14,000 globally** ([23] itbudgetcalculator.com).

Within that broader IT envelope, KPMG's global survey found the life sciences sector notably more conservative than cross-sector peers on digital technology overall: **97% of organizations allocate less than 1% of annual revenue to digital technology**, described as "behavior that matches the cross-sector average," and **93% report that digital initiatives generate less than 1% of annual revenue** in return ([24] assets.kpmg.com) ([25] assets.kpmg.com). This apparent contradiction, high reported AI integration alongside low digital-technology revenue share, reflects the fact that AI initiatives are often funded from existing R&D, IT, or commercial budgets rather than as a separately tracked line item, which is precisely why R&D-tech-budget-specific benchmarks are more useful for planning than total-revenue figures.

Benchling's 2026 Biotech AI Report, based on a November 2025 survey of approximately **100 biotechnology and pharmaceutical organizations actively using AI** (59% biopharma, 41% biotech; 53% with over 1,000 employees, 47% under), provides the clearest R&D-specific benchmark available. Respondents reported the following distribution of R&D tech budget allocated to AI: **less than 5% (27% of respondents)**, **5 to 10% (28%)**, **11 to 20% (28%)**, **21 to 35% (18%)**, and **more than 35% (6%)** ([26] www.benchling.com). In practice, this means the median actively-AI-using biotech or pharma organization allocates somewhere between 5% and 20% of its R&D technology budget to AI, with a meaningful long tail spending over one-fifth. It is worth noting this survey specifically targeted organizations already using AI, so the distribution likely overstates allocation intensity relative to the sector as a whole, which KPMG's broader, less AI-selected sample suggests remains far more conservative.

For mid-market planning at the organizational-function level rather than the percentage-of-revenue level, Sakara Digital's analysis of AI Center of Excellence structures for mid-cap pharma (defined as companies with **$500 million to $5 billion in revenue**) recommends a realistic annual budget of **$2 million**, against a large-cap enterprise AI investment range of **$50 million to $200 million annually** ([27] sakaradigital.com), noting mid-cap pharma more broadly operates in a **$5 million to $20 million** annual AI envelope across the whole enterprise, of which the CoE captures a focused slice. *Table 1* below summarizes the recommended $2 million CoE allocation.

Table 1 below summarizes the recommended annual budget allocation for a $2 million mid-cap pharma AI Center of Excellence.

Category	Annual Allocation	Share	What It Covers
Core staffing	$1.2M	60%	4 to 5 full-time equivalent (FTE) staff at competitive but not premium compensation
Technology and tooling	$400K	20%	Core platforms, model access, observability, governance tooling
External partnerships	$250K	12.5%	Strategic consulting, specialized capability buy-in for first deployments
Training and development	$100K	5%	Internal AI literacy, technical upskilling, conference attendance
Experimentation reserve	$50K	2.5%	Pilot funding, proof-of-concept work, evaluation licenses

Source: Sakara Digital, "Building an AI Center of Excellence on a $2M Budget: A Mid-Cap Pharma Approach" ([28] sakaradigital.com).

The allocation is deliberately conservative on technology (20%, versus what larger CoEs often spend) on the logic that mid-cap organizations should buy access to AI capabilities through Software-as-a-Service (SaaS) rather than build custom infrastructure, per the same source cited in Table 1 above. Compensation for core staffing in this model runs **$240,000 to $300,000 in loaded cost per FTE on average**, described as competitive for mid-cap pharma markets but not aspirational for top-tier AI talent. At the broader enterprise level (not specific to biotech), CFO-facing guidance suggests mid-market companies overall budget between **2% and 5% of revenue for AI**, with most clustering at 2.5 to 3.5%, and with that spend breaking down as roughly 40% tools and infrastructure, 30% talent, 20% integration, and 10% governance ([29] aisavvy.io).

AI Implementation Timelines and Adoption Stages

Mid-market biotech leaders planning an AI initiative need a realistic timeline, not the accelerated narrative that headline case studies sometimes imply. The data suggest a consistent multi-stage pattern: initial pilots move relatively fast, but scaling from pilot to enterprise-wide production remains the primary bottleneck across the sector.

ICON's 2024 survey (via GlobalData) of roughly 100 pharma and biotech research and development professionals in North America and Europe found that **49% of organizations are currently using AI or Big Data in their research programs**, up from about 39% in 2019, but only **13% report a comprehensive AI program fully implemented**, with roughly 70% still at pilot or selective-use stages (intuitionlabs.ai). BioPharm International's October 2025 digital transformation survey similarly found that while nearly **80% of respondents see AI and advanced analytics as the top investment area for the next three to five years**, **68% remain in pilot or exploratory phases** for AI specifically ([31] intuitionlabs.ai) ([32] intuitionlabs.ai).

ZS's 2026 survey of 115 U.S.-based pharma and biotech technology executives offers the most granular breakdown of where value materializes on which timeline. Roughly half of respondents (**49%**) say they are already consistently demonstrating measurable value through enterprise technology and data operations, and **47%** say the same for commercial sales and marketing, both considered "fast track" use cases ([33] www.zs.com). Discovery and clinical work follow a much longer "path two": only **17% of respondents can prove measurable value today in discovery**, though **42% expect value within the next year**; clinical shows a similar pattern, with about **30% demonstrating value now and 45% anticipating progress within the year** ([34] www.zs.com). Critically, the same survey found that **only about 40% of AI pilots that get initial funding eventually reach scaled deployment** ([35] www.zs.com), and **68% of respondents say neglecting data quality and governance early is the main reason AI initiatives fail**.

At the industry-wide level, Deloitte's 2026 Life Sciences Outlook Survey confirms this scaling gap is not improving as fast as adoption enthusiasm: **48% of respondents identified accelerated digital transformation** as a trend expected to substantially impact their organization in 2026, and **41% cited generative AI proliferation specifically**, yet **only 22% of life sciences leaders said they have successfully scaled AI, and just 9% reported achieving significant returns** on those efforts ([36] www.deloitte.com) ([37] www.deloitte.com). Slalom's separate research corroborates the gap with its own figures: **92% of life sciences leaders plan to increase AI investment, but only 33% can quantify its value** ([38] www.slalom.com), and **62% still run more than half their core applications on legacy platforms**, which limits interoperability and constrains how fast pilots can scale ([39] www.slalom.com).

For the mid-cap AI Center of Excellence model specifically, Sakara Digital recommends an explicit 18-month roadmap: months 1 to 3 for foundation-building (staffing, governance, AI inventory), months 4 to 9 for first deployments (executing two to three signature initiatives), months 10 to 15 for capability expansion (platform hardening, pattern transfer), and months 16 to 18 for a strategic reset ahead of the year-two budget decision ([40] sakaradigital.com). This roadmap implies that mid-market organizations should not expect production-grade, validated AI deployments inside their regulated functions in under roughly six months from a standing start, and should expect the second-year budget decision, not the first, to be the real test of whether the investment scales.

Infrastructure, Talent, and Vendor Cost Drivers

Three cost categories dominate a biotech AI budget: compute infrastructure, specialized talent, and software or platform licensing. Each has moved meaningfully in 2025 and 2026, generally in favor of buyers, even as headline enterprise AI budgets keep expanding.

Compute infrastructure. Graphics processing unit (GPU) cloud pricing, the primary driver of AI training and inference cost for computational drug discovery, has fallen sharply. AWS cut its P5 H100 instance list pricing by **44% in June 2025**, bringing on-demand cost to roughly **$6.88 per GPU-hour** (an 8-GPU p5.48xlarge instance costs $55.04 per hour) (www.spheron.network). Alternative "neocloud" providers now offer the same NVIDIA H100 hardware from **under $2 per GPU-hour on-demand**, with providers such as Lambda Labs listing roughly $2.49 per GPU-hour and RunPod roughly $3.29 per GPU-hour (www.spheron.network), and Spheron's own spot pricing starting as low as **$1.66 per GPU-hour**, versus AWS's rarely-available P5 spot rate of roughly $3.83 per GPU-hour (www.spheron.network). For a persistent 8-GPU training cluster running full time (720 hours per month), the gap between AWS's 3-year reserved pricing ($17,834 per month) and a neocloud on-demand rate ($15,206 per month) illustrates that reserved hyperscaler commitments do not automatically beat flexible alternatives on pure compute cost (www.spheron.network). At the enterprise scale end of the spectrum, Roche's newly expanded on-premise and cloud NVIDIA infrastructure, anchored by an additional **2,176 Blackwell GPUs** announced in March 2026, now totals **more than 3,500 GPUs**, described as the largest disclosed GPU footprint available to a pharmaceutical company ([41] www.roche.com) ([42] www.roche.com). Separately, Eli Lilly and NVIDIA announced in January 2026 that they will **"jointly invest up to $1 billion over five years"** in talent, infrastructure, and compute for a co-innovation AI lab combining Lilly's drug-development expertise with NVIDIA's accelerated computing platform ([43] nvidianews.nvidia.com) ([44] nvidianews.nvidia.com). These figures make clear that a mid-market biotech company's realistic infrastructure envelope is orders of magnitude smaller than large-cap peers, reinforcing why the mid-cap CoE model described above recommends buying commercial application programming interface (API) access to foundation models rather than pursuing private GPU deployment, which becomes economic only at a far higher utilization level than most mid-cap organizations reach in their first two years of an AI program.

Talent. Specialized AI engineering talent remains the largest single cost line in most biotech AI budgets. Average base salaries for AI engineers in the United States reached roughly **$206,000 in 2025**, a **$50,000 increase** over the prior annual cycle, according to salary benchmarking firm Signify Technology ([45] www.signifytechnology.com), and 2026 AI engineer base salaries broadly span **$145,000 to $310,000** depending on seniority and location ([46] www.kore1.com). Benchling's survey found biotech's dominant source of AI talent is not external hiring at all: **67% of respondents cite internal upskilling of existing scientific staff** as the primary talent source, versus only **21% from tech companies** and **19% from external contractors and consulting partnerships** ([47] www.benchling.com). This matters directly for mid-market budgeting because it suggests the more capital-efficient path is training existing bench scientists into "trilingual" science-plus-AI roles rather than competing for scarce, expensive AI-native engineering hires against large-cap pharma and well-funded AI-native biotechs, a talent market disadvantage industry advisors explicitly flag as a structural mid-cap constraint ([48] sakaradigital.com).

Software and platform licensing. For enterprise life sciences software that underpins much AI-enabled workflow deployment, contract data compiled by procurement-benchmarking firm Vendr shows a **median annual contract value of $211,872**, with a range of **$113,755 to $501,612**, for Veeva Systems' cloud platform products, which do not publish public pricing and negotiate every enterprise contract individually ([49] checkthat.ai). Vendr's own live marketplace page, checked in this report, currently lists a somewhat lower **median contract value of $188,460**, with a range of **$90,720 to $491,710** ([50] www.vendr.com), a discrepancy that itself illustrates a broader point: negotiated enterprise software benchmarks shift over time and by sample, and mid-market buyers should treat any single published contract-value figure as directional rather than precise. More granular estimates put Veeva customer relationship management (CRM) licensing at roughly **$150 to $250 per user per month** and Vault products at **$100 to $200 per user per month** ([51] ai-scanner.com). This pricing opacity itself has budgeting implications: a mid-market biotech evaluating a first enterprise AI-enabled platform deployment should expect a negotiated, not published, quote, and should budget procurement and negotiation time accordingly rather than assuming a self-service SaaS pricing model. ZS's CDIO survey confirms this infrastructure-first spending pattern extends well beyond software licensing: **88% of surveyed pharma and biotech technology executives are increasing investment in cloud and infrastructure** over the next 12 months, followed by data products and platforms (**86%**), AI platforms specifically (**84%**), and IT and DevOps tooling (**79%**) ([52] www.zs.com). Security budgeting deserves its own explicit line: industry best practice allocates **10 to 15% of total IT budget to cybersecurity** at most organizations, rising to **15 to 18% for regulated industries** including healthcare, with measured security spend averaging **10.9% of IT budget in 2025** ([53] itbudgetcalculator.com), a figure mid-market biotech AI budgets should fold into the technology and tooling line rather than treating as separate. Of that security allocation, roughly **40% typically goes to software and platforms, 30% to personnel, 15% to hardware, and 15% to outsourced services** ([54] itbudgetcalculator.com), a breakdown mid-cap CoE governance line items should mirror when they extend into AI-specific security controls.

Analysis of Key Segments: Budget Benchmarks by Company Stage

Biotech AI spending behaves differently depending on company stage, and mid-market leaders benchmarking against the wrong peer group will draw the wrong conclusions. Three distinct segments are visible in the assembled data: early-stage, venture-backed biotechs; mid-cap organizations with an approved product or late-stage pipeline; and large-cap incumbents.

Early-stage, venture-backed biotechs show the widest variance in AI-related infrastructure spend because compute cost, not headcount, often dominates their budget. Rezo, a growth-stage biotech running protein-protein interaction modeling with tools like AlphaFold, worked with orchestration vendor Union.ai to cut storage costs by **up to 90%** and compute costs by **up to 67%**, in one case reducing the monthly cost of multiple sequence alignment (MSA) workloads from roughly **$10,000 a month to hundreds of dollars** by shifting to on-the-fly local solid-state drive (SSD) provisioning ([55] www.union.ai) ([56] www.union.ai). Notably, Rezo achieved this without adding dedicated DevOps headcount, illustrating that small teams can access enterprise-scale compute elasticity, up to 500 tensor processing units (TPUs) across four regions, through orchestration platforms rather than in-house infrastructure engineering ([57] www.union.ai). This pattern, extreme compute elasticity purchased rather than built, recurs across other early-stage examples: Terray Therapeutics uses NVIDIA DGX Cloud to train foundation models for chemistry rather than building its own data center ([58] www.nvidia.com), and UK-based RNA interference (RNAi) biotech Tangram Therapeutics built its "LLibra OS" agentic AI platform on AWS, achieving a **300x increase in data processed daily** and a **50x acceleration in target hypothesis evaluation**, processing **over 1,000 datasets using AWS Glue**, without a comparable increase in fixed infrastructure spend ([59] aws.amazon.com) ([60] aws.amazon.com).

Mid-cap organizations, the primary audience for this report, sit between these venture-funded infrastructure-light models and large-cap capital intensity. As established in the Budget Benchmarks section, the $2 million AI Center of Excellence envelope recommended for $500 million to $5 billion revenue mid-cap pharma companies deliberately mirrors the early-stage pattern of buying rather than building compute and platform capability, while adding the staffing and governance layers that early-stage companies can defer but regulated mid-caps cannot. A regional example of mid-market-adjacent adoption comes from an EY-Parthenon India analysis of over 500 pharma roles, where **50% of firms report investing in or exploring AI-driven solutions** and **25% already have generative AI (GenAI) tools in production**, with the firm estimating **30 to 40% potential productivity gains** from GenAI in sales, supply chain, and production ([61] pharma.economictimes.indiatimes.com) ([62] pharma.economictimes.indiatimes.com). The same analysis found **75% of surveyed companies reported cost reductions and increased customer satisfaction** attributable to AI-driven tools ([63] pharma.economictimes.indiatimes.com), a pattern broadly consistent with the mid-cap “buy over build” thesis even outside the US and European markets covered by most other sources in this report.

Large-cap incumbents occupy the top of the spend distribution and set the reference points that mid-market leaders should explicitly avoid benchmarking against. Genentech’s generative AI research agent, built on Amazon Bedrock, is projected to automate **over 43,000 hours of manual effort annually** in biomarker validation, work that previously consumed close to five years of scientist time in aggregate across therapeutic areas, and compresses tasks that took weeks into minutes ([64] aws.amazon.com) ([65] aws.amazon.com). Large-cap organizations can absorb the multi-year, multi-hundred-million-dollar infrastructure and research investment such outcomes require; a mid-cap biotech generally cannot, and should instead look to the compute-elastic, platform-first patterns visible in the early-stage segment as the more transferable model.

A cautionary counter-signal across all three segments comes from a 2024 ZoomRx survey of over 200 life sciences professionals (via Fierce Pharma), where **83% of respondents called AI “overrated”** even though **92% said their firms are actually using AI in some capacity**, with 50% reporting “some” use cases in production and 10% describing themselves as industry-leading in AI adoption ([66] www.fiercepharma.com). The same survey found **64% of respondents cite cost savings as their primary AI motivation**, versus only **17% who see AI as having a significant impact on revenue** ([67] www.fiercepharma.com), and **65% of the top 20 global pharmaceutical companies have banned employee use of consumer generative AI tools like ChatGPT** over data leakage concerns, even as more than half of surveyed professionals use such tools regularly anyway ([68] www.fiercepharma.com). The same respondents show a wide trust gap even where they use AI daily: **81% expressed belief in AI’s ability to enhance efficiency and effectiveness**, yet **91% expressed concern about AI’s impact on data security, privacy, and use** ([69] www.fiercepharma.com). This gap between sentiment (skeptical) and behavior (widespread, informal use) is itself a budget risk: mid-market organizations that have not formally budgeted for AI governance are very likely already carrying uncontrolled AI-related data exposure through informal employee use of consumer tools, a hidden cost this report’s governance recommendations are designed to address.

ROI of AI in Drug Discovery: What the Evidence Supports

Return on investment (ROI) claims for AI in drug discovery span an unusually wide range, from vendor-reported efficiency multiples in the hundreds of percent to independent survey data showing only a small minority of organizations can currently prove measurable value. Mid-market leaders building a business case should draw on both categories of evidence and be explicit about which is which.

On the vendor and case-study side, reported efficiency gains are often dramatic at the task level. NVIDIA’s Innophore case study reports that a deep learning model reduced a single protein-binding prediction task from

2,000 CPU hours to 15 minutes, a **95% cost reduction**, and cut a related biological representation task from over a year to two weeks ([70] www.contextwindows.ai). Historically, Exscientia's collaboration with Sumitomo Dainippon Pharma produced the first AI-designed small molecule (DSP-1181) to enter Phase 1 clinical trials, reportedly completed in under 12 months against a traditional design timeline of five or more years, though the compound later failed in trials, a reminder that discovery-speed gains do not guarantee clinical success (www.itmedia.co.jp) ([71] intuitionlabs.ai). More recently, Recursion identified its CDK7 inhibitor candidate REC-617 as a lead candidate in **under 11 months**, synthesizing **136 novel compounds** using its AI-driven precision design platform, before advancing it into combination Phase 1/2 studies ([72] ir.recursion.com). Independent industry analysis places broader AI-attributable timeline compression at the preclinical stage specifically, with AI integration able to **shorten preclinical stages from roughly six years to two and a half years** in some programs ([73] www.slalom.com).

On the survey side, the picture is more measured. Benchling's 2026 survey of active AI users found **50% of biotech report faster time-to-target today**, and **56% expect cost reductions within two years** as automation and agentic workflows scale further, with the clearest AI wins concentrated in specific, well-bounded workflows: literature and knowledge extraction (**76% adoption**), protein structure and property prediction (**71%**), and scientific reporting (**66%**), while adoption drops sharply in less bounded, more judgment-intensive areas like biomarker analysis (**40%**) and generative de novo drug design (**42%**), where underlying data is scattered, incomplete, and harder to validate ([74] www.benchling.com). KPMG's broader survey found **75% of life sciences technology leaders trust AI outputs for strategic and operational decisions**, and nearly half (**48%**) report high-tier digital technology returns exceeding **$100 million**, with the largest concentration in the $250 million to $499.9 million range ([75] assets.kpmg.com) ([76] assets.kpmg.com). Critically, however, this financial value is concentrated among large-cap respondents whose absolute technology budgets can produce nine-figure returns; a mid-cap organization working from a $2 million to $20 million annual AI envelope should expect proportionally smaller absolute returns even at a comparable percentage ROI, and should build its business case around percentage-based operational metrics (cycle-time reduction, compounds synthesized per program, hours automated) rather than absolute dollar-value claims lifted from large-cap case studies.

The most disciplined summary of ROI realism comes from ZS's CDIO survey, where pharma chief information and digital officers reported concentrating investment on **five to 10 high-impact use cases**, each targeted at **20% to 30% ROI**, rather than spreading budget across a broad portfolio of speculative pilots ([77] www.zs.com). This concentrated, use-case-specific approach to ROI measurement, rather than a single blended "AI ROI" figure for the whole organization, is the model this report recommends mid-market biotech leaders adopt when presenting a business case to their board or investment committee.

Data readiness is the precondition behind most of these ROI figures, and it is where high-AI-adoption biotechs measurably differ from low-adoption peers. The same Benchling survey referenced above found that among biotechs with high AI adoption, **33% report "advanced" data integration** (most datasets standardized and accessible with reliable pipelines), compared with only **25% among low-adoption biotechs**. Wet-dry lab integration shows a similar gap: **30% of high-AI-adoption biotechs report highly integrated wet-dry lab workflows**, versus **18% of low-adoption biotechs**, per the same survey. This correlation, not proof of causation, is nonetheless consistent with the ZS finding that 68% of respondents cite data quality and governance failures as the primary reason AI initiatives fail, and it reinforces why this report's recommended mid-market strategy treats data infrastructure as a prerequisite investment rather than a parallel workstream.

Building a Mid-Market AI Enablement Strategy

Translating the budget, timeline, and ROI benchmarks above into an executable strategy requires mid-market biotech leaders to make a small number of explicit, sequenced decisions rather than launching parallel, uncoordinated pilots. Several patterns recur across the advisory and survey sources in this report:

- **Anchor the budget to a denominator before naming a number.** Decide whether a given initiative is R&D-facing (benchmark against the 5 to 20% R&D-tech-budget range established above) or enterprise-wide (benchmark against the 6 to 9% healthcare/life sciences IT-spend-as-percentage-of-revenue range) ([4] itbudgetcalculator.com), and do not blend the two.
- **Staff before you buy infrastructure.** The $2 million mid-cap CoE model outlined in Table 1 allocates 60% to staffing and only 20% to technology, and survey data shows internal upskilling, not external hiring, is the dominant AI talent source in biotech today (67% of respondents, as detailed in the Infrastructure, Talent, and Vendor Cost Drivers section above).
- **Buy commercial foundation-model access; do not build a proprietary model.** This is the dominant sourcing pattern across surveyed organizations, per the survey data cited above: 60% buy commercial AI applications and 56% fine-tune third-party models with proprietary data, versus only 48% training and deploying fully first-party models internally. This mirrors KPMG's broader finding that a hybrid build-and-buy approach is the preferred capacity-building model for **45% of life sciences organizations**, with pure external buying the least-favored option at just **15%** ([78] assets.kpmg.com).
- **Fund governance as its own line item, not an afterthought.** The January 2026 FDA-EMA joint guiding principles formalize documentation, risk-based assessment, and life-cycle management obligations that mid-market organizations should budget for explicitly, ideally through a dedicated AI quality and governance role rather than splitting the work informally across existing staff ([79] www.fda.gov).
- **Prioritize a small number of measurable use cases over broad experimentation.** The finding that CIOs are concentrating on five to 10 use cases targeted at 20 to 30% ROI, as discussed in the ROI section above, rather than a broad pilot portfolio, is the pattern most consistent with the sector's low pilot-to-scale conversion rate.
- **Set an explicit second-year budget checkpoint.** Rather than treating year one as sufficient proof, structure the initiative so that year-two funding is contingent on documented value, consistent with the mid-cap CoE roadmap's emphasis on the economics of the second-year decision described above.

These six decisions do not guarantee AI success, but they align a mid-market biotech's spending pattern with what the assembled 2026 benchmark data shows actually correlates with organizations that report measurable value rather than stalled pilots.

Data Analysis and Evidence

The quantitative backdrop against which mid-market biotech AI budgets should be benchmarked is a market growing extremely fast in aggregate while remaining highly concentrated at the top of the market. Worldwide AI spending across all industries is forecast to total **$2.52 trillion in 2026**, a **44% increase year-over-year**, driven substantially by a **49% increase in spending on AI-optimized servers** and **$401 billion in incremental AI infrastructure spending** as technology providers build out AI foundations ([80] www.gartner.com) ([81] www.gartner.com). Broader worldwide IT spending, of which AI is a growing subset, reaches **$6.15 trillion** in 2026, with data center systems spending alone rising **31.7%** to surpass **$650 billion**, up from nearly $500 billion the year before ([82] www.gartner.com). Generative AI (GenAI) software specifically is forecast to grow **80.8%** in 2026, and worldwide software spending overall is projected to remain above **$1.4 trillion**, a growth rate that Gartner attributes partly to GenAI's rising share of the software market ([83] www.gartner.com).

NVIDIA's own 2026 state-of-AI survey work, spanning industries, found that **86% of respondents said their AI budget will increase in 2026**, with another **12% saying budgets will stay flat** and roughly **40% expecting increases of 10% or more**; healthcare and life sciences was specifically named among the industries "showing the strongest adoption and ROI results" ([84] blogs.nvidia.com) ([85] blogs.nvidia.com). Within that spend, **42% of**

respondents named optimizing AI workflows and production cycles as the top 2026 spending priority, ahead of finding new use cases (31%) and building AI infrastructure (31%) ([86] blogs.nvidia.com). This is a meaningful signal for mid-market biotech budgeters: the dominant 2026 industry spending pattern is optimization of existing deployments, not net-new infrastructure buildout, which supports a "buy platform, build differentiation" posture over large capital outlays on custom infrastructure. Across all industries and regions surveyed, **64% of respondents said their organizations are actively using AI in operations**, with 28% still in the assessment phase and just 8% not using AI at all ([87] blogs.nvidia.com). Regionally, NVIDIA's cross-industry survey found North America leads AI adoption at **70% actively using the technology**, with large companies (over 1,000 employees) showing markedly broader adoption than smaller peers: **76% of large-company respondents report active AI usage**, versus lower rates at smaller organizations, a gap the survey attributes to large companies having "more capital to invest in AI infrastructure, data scientists and experts" ([88] blogs.nvidia.com) ([89] blogs.nvidia.com). For mid-market biotech leaders, this capital-scale adoption gap is precisely the dynamic the budget benchmarks in this report are designed to help correct: a smaller absolute budget, deployed against the buy-over-build, staffing-first framework outlined above, can still close a meaningful share of that adoption gap. On the budget-commitment side, NVIDIA's survey found **nearly all respondents (86%) said their AI budget will increase in 2026**, with North American organizations especially aggressive: **48% of North American respondents said their AI budget would increase by 10% or more**, as did 45% of executive-level respondents specifically ([90] blogs.nvidia.com).

Venture and deal-flow data confirms life sciences AI is attracting a disproportionate share of available capital even as overall healthcare investment contracts. Silicon Valley Bank's 2026 Healthcare Investments and Exits report found AI investment in healthcare **exceeded $18 billion in 2025**, representing **46% of all healthcare venture investment**, even as total healthcare investment fell **12%** to **$46.8 billion** from a 2021 peak of $68.3 billion ([91] www.svb.com) ([92] www.prnewswire.com). Notably, deals over $300 million now account for **40% of total healthcare AI spending**, up from 31% in 2024 and 29% in 2023, indicating capital is concentrating into fewer, larger AI-native platform bets rather than spreading broadly ([93] www.prnewswire.com). This concentration effect is a structural headwind mid-market biotechs should factor into fundraising and partnership strategy: capital increasingly favors either AI-native platform companies or established large-cap incumbents, leaving traditional mid-cap biotechs to compete on capital-efficient execution rather than infrastructure scale.

Table 2 below summarizes the survey-based AI budget allocation benchmarks gathered across the primary sources in this report, since methodologies and sample populations differ meaningfully and should not be treated as interchangeable.

Source	Sample	Metric	Key Finding
KPMG Global Tech Report 2026	124 Life Sciences tech leaders	Digital tech spend as % of revenue	97% spend under 1% of revenue on digital technology ([94] assets.kpmg.com)
Benchling 2026 Biotech AI Report	~100 AI-using biotech/pharma orgs	% of R&D tech budget on AI	52% allocate 11% or more; 27% allocate under 5% ([95] www.benchling.com)
Sakara Digital / mid-cap pharma	Advisory engagements, mid-cap pharma	Annual AI CoE budget	$2M realistic CoE envelope; $5M to $20M total enterprise AI spend (see Table 1 above)
itbudgetcalculator.com (Gartner/Avasant-sourced)	Cross-industry benchmark	IT spend as % of revenue	Healthcare/life sciences 6 to 9%, versus 5.7% cross-industry average ([4] itbudgetcalculator.com)
AI Savvy (CFO-grade guidance)	Mid-market companies, all sectors	AI budget as % of revenue	2 to 5% of revenue, most at 2.5 to 3.5% ([5] aisavvy.io)

The spread across these benchmarks (from under 1% of revenue at the conservative KPMG end to over 20% of R&D tech budget at the aggressive Benchling end) is not a contradiction so much as a reflection of denominator choice: total revenue, total IT budget, or R&D technology budget specifically produce very different percentages for the same absolute dollar figure. Mid-market biotech leaders should anchor their own budget conversation to the R&D-tech-budget denominator when the AI initiative is R&D-facing, and to the general IT-budget denominator when it spans commercial, manufacturing, or enterprise functions, rather than defaulting to a single blended percentage that obscures which spending pool is actually being discussed.

Case Studies and Real-World Examples

Pfizer and AWS: Infrastructure Cost Reduction Alongside R&D Acceleration

Pfizer's collaboration with AWS under the Pfizer-Amazon Collaboration Team (PACT) initiative, spanning 14 projects including generative AI and machine learning applications, delivered two quantified outcomes: **16,000 hours of scientist search time saved annually**, and a **55% reduction in infrastructure costs** ([96] aws.amazon.com). Separately, Pfizer's use of supercomputing and AI in its COVID-19 antiviral program (Paxlovid) is credited with reducing a key manufacturing cycle by **67%**, allowing the company to produce thousands more doses per batch, and the company now uses AI in over **half of its clinical trials** ([97] intuitionlabs.ai) ([98] intuitionlabs.ai). For mid-market biotech leaders, the relevant lesson is not the absolute scale (Pfizer is a large-cap company with resources far beyond a mid-market peer) but the pattern: infrastructure cost reduction and R&D acceleration were pursued as parallel, quantified objectives from the outset, rather than treating cost savings as an afterthought to a discovery-focused pilot.

Roche and NVIDIA: The Scale Ceiling for Compute-Intensive AI

Roche's March 2026 expansion of its NVIDIA-powered AI infrastructure added **2,176 Blackwell GPUs** on-premise across the United States and Europe, bringing its combined on-premise and cloud GPU footprint to **more than 3,500 GPUs**, which the company describes as the greatest disclosed GPU footprint available to any pharmaceutical company ([99] www.roche.com). This deployment, embedded across Roche's value chain to accelerate both diagnostics and therapeutics development, illustrates the ceiling of private compute investment that only the largest life sciences companies can justify. It is a useful negative benchmark for mid-market planning: a company with revenues in the low single-digit billions has no realistic path to matching this infrastructure scale, which is precisely why the mid-cap CoE framework described earlier in this report recommends commercial API access over private GPU deployment for the large majority of mid-cap organizations.

Recursion Pharmaceuticals: A Clinical-Stage "TechBio" Cost Structure

Recursion Pharmaceuticals, formed through a 2024 combination with Exscientia, reported full-year 2025 **cash operating expense of approximately $400 million**, roughly **10% lower than its original guidance**, with 2026 cash operating expense guided to **under $390 million** ([100] ir.recursion.com) ([101] ir.recursion.com). The company ended 2025 with **$753.9 million in cash and cash equivalents**, providing runway into early 2028

([102] ir.recursion.com). Operationally, Recursion reports its AI-native platform (the "Recursion OS") delivered development candidates by synthesizing roughly **330 compounds per program over approximately 17 months**, compared with an industry average of over **2,500 compounds and 42 months** (ir.recursion.com). The company also reports having generated **over $500 million in upfront and progress-based milestone payments** from partnerships with Sanofi (**$134 million** to date across five accepted discovery program packages, with the collaboration structured for up to **15 AI-designed small molecule programs** in total) and Roche and Genentech (**$213 million** to date) ([103] ir.recursion.com) ([104] ir.recursion.com) ([105] ir.recursion.com). For mid-market biotech leaders, Recursion's structure (roughly $400 million in annual operating cost against a differentiated AI-native platform) illustrates the far end of "AI-first" biotech spend, well beyond what a traditional mid-cap organization retrofitting AI into existing operations would need, but the compounds-per-program and months-per-program metrics offer a useful external benchmark for evaluating whether a smaller-scale AI discovery deployment is actually accelerating candidate generation.

Insilico Medicine: Funding Trajectory of an AI-Native Discovery Platform

Insilico Medicine's funding history illustrates the capital intensity required to build a proprietary AI drug discovery platform from the ground up: the company raised a total of **$510 million across 12 funding rounds** since founding, including a **$255 million Series C** and a subsequent Series D and E, before its Hong Kong initial public offering (IPO) brought cash and bank balances to **$393.3 million** as of the most recent reporting ([106] exa.ai) ([107] insilico.com). The company's 2023 collaboration with Eli Lilly, structured as a research and licensing agreement, was worth up to **$2.75 billion** in potential milestones ([108] www.pharmexec.com). This case underscores that building a fully proprietary AI discovery platform, as opposed to licensing or partnering for AI capability, is a venture-capital-scale undertaking measured in hundreds of millions of dollars, not a line item inside an existing mid-cap R&D budget, reinforcing the buy-versus-build calculus mid-market organizations must make explicitly rather than by default.

FDA and EMA: Regulatory Guardrails Shaping Implementation Cost and Timeline

On **January 14, 2026**, the FDA's Center for Drug Evaluation and Research (CDER) and Center for Biologics Evaluation and Research (CBER) jointly published, in collaboration with the EMA, **10 guiding principles for "Good AI Practice" in drug and biological product development** ([109] www.fda.gov) (www.ema.europa.eu). The principles are described as "a first step of a renewed EU-US cooperation in the field of novel medical technologies," according to European Commissioner for Health and Animal Welfare Olivér Várhelyi, and will underpin future AI-specific guidance across both jurisdictions (www.ema.europa.eu). The principles cover human-centric design, a risk-based approach, adherence to standards, clear context of use, multidisciplinary expertise, data governance and documentation, model design and development practices, risk-based performance assessment, life cycle management, and clear essential information ([110] www.fda.gov). This is a Tier 1, non-hypothetical regulatory case relevant to budget planning because it formalizes documentation, validation, and governance obligations that were previously informal best practice, meaning mid-market biotech AI budgets that omit a dedicated governance and validation cost line, of the kind captured in a dedicated "AI quality and governance" staffing role, are now underbudgeting against an explicit regulatory expectation rather than merely an internal best practice.

Implications and Future Directions

The convergence of the data assembled in this report points toward several practical implications for mid-market biotech leaders building or revising an AI budget in the second half of 2026. First, the denominator problem described in the Data Analysis section is not a data-quality artifact to be resolved but a structural feature of how AI spend gets tracked, and mid-market organizations should adopt at least two parallel benchmarks (R&D-tech-budget share for R&D-facing initiatives, general IT-budget share for enterprise-wide initiatives) rather than seeking a single blended percentage that will inevitably mislead one stakeholder group or another.

Second, the gap between adoption breadth and scaling depth, visible consistently across KPMG (44% report limited AI scaling maturity), Deloitte (only 22% have successfully scaled AI, 9% report significant returns), and ZS (only 40% of funded pilots reach scaled deployment), should reset mid-market planning expectations away from a single-year “prove AI works” pilot toward a multi-year, staged investment with an explicit second-year go/no-go decision point, consistent with the Sakara Digital 18-month roadmap and its emphasis on the year-two budget conversation as the true test of program viability. Organizations that treat the first 12 months as sufficient to prove enterprise-wide value are working against the grain of what the sector-wide data actually shows.

Third, falling GPU compute costs (a 44% AWS price cut in mid-2025 alone, with neocloud alternatives pricing meaningfully below even the reduced hyperscaler rate) mean the marginal cost of computational experimentation is dropping faster than the marginal cost of the talent and governance functions required to use that compute responsibly in a regulated environment. This shifts the center of gravity in a mid-cap AI budget away from infrastructure and toward people and process, a trend the $2 million mid-cap CoE allocation (60% staffing, only 20% technology) already reflects and which is likely to become more pronounced, not less, as compute commoditizes further.

Fourth, the January 2026 FDA-EMA joint guiding principles are an early but consequential signal that regulatory expectations for AI governance in drug development are converging internationally and formalizing quickly, a pattern Deloitte’s own survey corroborates: **51% of non-US life sciences respondents cite national regulatory changes, including the EU AI Act, as a trend expected to shape their organizational strategy in 2026** ([111] www.deloitte.com). Mid-market biotech organizations that have historically treated AI governance as an internal, informal function should expect this cost category, documentation, validation evidence, and risk-based performance assessment, to grow as a share of total AI budget over the next several budget cycles, not shrink, even as raw compute and even talent costs potentially ease.

Finally, capital concentration in AI-native platform companies (evidenced by SVB’s finding that deals over $300 million now represent 40% of healthcare AI investment, up from 29% two years prior) suggests the venture funding environment increasingly rewards AI-native business models over AI-retrofitted incumbents. For a traditional mid-cap biotech, this argues for a partnership-and-licensing-forward AI strategy, working with platform companies like Recursion or Insilico, or with cloud/infrastructure partners like AWS and NVIDIA, over attempting to build a fully proprietary discovery-AI stack that only venture-scale capital can sustain.

A sixth implication concerns workforce readiness, which the survey data suggests is now a larger constraint on AI ROI than either budget size or technology availability. Slalom’s research found **56% of life sciences leaders cite workforce skill gaps as the top barrier to AI return on investment**, ahead of governance or technology constraints ([112] www.slalom.com), which reinforces this report’s recommendation that mid-cap budgets weight internal upskilling over external hiring. Separately, cybersecurity concerns are rising in parallel with AI adoption: **35% of Deloitte’s 2026 survey respondents said cybersecurity will affect their organizational strategy** this year, up from prior years, a trend mid-market AI budgets should account for given that AI systems widen the data-access surface that security teams must defend ([113] www.deloitte.com).

Frequently Asked Questions (FAQs)

What percentage of revenue should a biotech company budget for AI? There is no single correct figure, but the assembled benchmarks suggest 2 to 5% of total revenue for AI specifically at mid-market companies generally ([5] aisavvy.io), while life sciences organizations broadly report allocating far less to digital technology overall (under 1% for 97% of KPMG's respondents) ([94] assets.kpmg.com), suggesting AI spend at most biotechs is currently funded from existing functional budgets rather than as a distinct revenue-percentage line item.

How long does it take to implement AI in a biotech R&D function? Initial pilots can produce measurable results within months in fast-track areas like enterprise data operations and commercial analytics, where roughly half of surveyed organizations already report consistent value, per the ZS survey data discussed in the Implementation Timelines section above. In discovery and clinical R&D specifically, only 17 to 30% of organizations report measurable value today, with most expecting meaningful results only within the following 12 months, and full enterprise scaling, per the mid-cap CoE roadmap discussed above, realistically spans 18 months from a standing start.

What is the biggest cost driver in a biotech AI implementation? Staffing and talent, not compute infrastructure, represents the largest cost category in the recommended mid-cap CoE budget structure detailed in Table 1 (60% of a $2 million annual envelope), a pattern reinforced by falling GPU cloud pricing that has reduced the relative cost of raw compute.

Is buying AI capability more cost-effective than building it in-house for a mid-market biotech? For the large majority of mid-cap organizations, yes: the mid-cap CoE framework discussed in the Budget Benchmarks section explicitly recommends against building custom foundation models, since the economics do not work at mid-cap scale, and instead recommends commercial API access to foundation models combined with SaaS tooling. Fully proprietary AI-native discovery platforms, as built by Recursion or Insilico Medicine, required hundreds of millions of dollars in venture and public capital to reach maturity ([114] exa.ai).

What regulatory requirements affect AI budget planning in biotech? The FDA and EMA's January 2026 joint "10 Guiding Principles" for Good AI Practice in drug development formalize expectations around data governance, model documentation, risk-based performance assessment, and life-cycle management ([79] www.fda.gov), meaning mid-market budgets should include a dedicated governance and validation function rather than treating compliance as incidental to a discovery or commercial AI pilot.

How much does cloud GPU compute cost for AI drug discovery in 2026? On-demand H100 GPU pricing on AWS runs approximately $6.88 per GPU-hour following a 44% mid-2025 price cut, while alternative cloud providers offer comparable hardware from under $2 per GPU-hour on-demand (www.spheron.network), meaning provider selection alone can shift a sustained multi-GPU training budget by 50% or more.

Conclusion

Biotech AI budget benchmarking in 2026 does not yield a single defensible number, and any consultancy or vendor offering one should be treated skeptically. What the assembled evidence supports instead is a structured planning framework: total digital technology spend for most life sciences organizations remains under 1% of revenue, R&D-tech-budget allocation to AI specifically clusters between 5% and 20% among organizations already using AI, and a realistic mid-cap pharma AI Center of Excellence operates on roughly $2 million annually, an order of magnitude below the $50 million to $200 million large-cap enterprise range. Implementation timelines remain longer than the AI industry's general narrative implies, with only a fifth of life sciences organizations reporting successful scaling and fewer still reporting significant returns, even as compute costs fall and adoption breadths widen. For a mid-market biotech leader building the business case for

AI investment, the actionable conclusion is to budget primarily for staffing, governance, and data foundations rather than compute infrastructure, to plan on an 18-month horizon before expecting a defensible scaling decision, and to treat regulatory guidance from the FDA and EMA as a budget line item rather than an afterthought. None of these figures should be treated as fixed; they represent the best available synthesis of survey, financial, and case-study data as of July 2026, and mid-market organizations should revisit them at least annually as the underlying AI cost and adoption curves continue to shift.

External Sources

[1] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:of%20...

[2] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:have%...

[3] https://www.benchling.com/biotech-ai-report-2026#:~:28%25...

[4] https://itbudgetcalculator.com/#:~:Healt...

[5] https://aisavvy.io/insights/how-much-should-we-budget-for-ai-this-year#:~:Mid-m...

[6] https://sakaradigital.com/blog/ai-center-of-excellence-2m-budget-mid-cap-pharma/#:~:is%20...

[7] https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026#:~:World...

[8] https://www.gartner.com/en/newsroom/press-releases/2026-02-03-gartner-forecasts-worldwide-it-spending-to-grow-10-point-8-percent-in-2026-totaling-6-point-15-trillion-dollars#:~:World...

[9] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Despi...

[10] https://aws.amazon.com/solutions/case-studies/pfizer-PACT-case-study/#:~:16%2C...

[11] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:this%...

[12] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:Advan...

[13] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:97%25...

[14] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:48%20...

[15] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:44%20...

[16] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:more%...

[17] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:35%20...

[18] https://docs.oracle.com/en/programs/research/hpc-and-drug-discovery-whitepaper/F59484_01.pdf#:~:typic...

[19] https://www.drugpatentwatch.com/blog/artificial-intelligence-in-drug-discovery-what-is-realistic-what-are-illusions-part-1-ways-to-make-an-impact-and-why-we-are-not-there-yet/#:~:rough...

[20] https://intuitionlabs.ai/articles/ai-adoption-pharma-biotech-benchmarks#:~:These...

[21] https://itbudgetcalculator.com/#:~:Finan...

[22] https://itbudgetcalculator.com/#:~:Mid-m...

[23] https://itbudgetcalculator.com/#:~:rough...

[24] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:The%2...

[25] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:93%20...

[26] https://www.benchling.com/biotech-ai-report-2026#:~:20%25...

[27] https://sakaradigital.com/blog/ai-center-of-excellence-2m-budget-mid-cap-pharma/#:~:A%20l...

[28] https://sakaradigital.com/blog/ai-center-of-excellence-2m-budget-mid-cap-pharma/#:~:Core%...

[29] https://aisavvy.io/insights/how-much-should-we-budget-for-ai-this-year#:~:Of%20...

[30] https://intuitionlabs.ai/articles/ai-adoption-pharma-biotech-benchmarks#:~:only%...

[31] https://intuitionlabs.ai/articles/ai-adoption-pharma-biotech-benchmarks#:~:Nearl...

[32] https://intuitionlabs.ai/articles/ai-adoption-pharma-biotech-benchmarks#:~:Yet%2...

[33] https://www.zs.com/insights/scaling-ai-in-pharma-cdio-2026#:~:Nearl...

[34] https://www.zs.com/insights/scaling-ai-in-pharma-cdio-2026#:~:In%20...

[35] https://www.zs.com/insights/scaling-ai-in-pharma-cdio-2026#:~:just%...

[36] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Nearl...

[37] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Despi...

[38] https://www.slalom.com/us/en/insights/life-sciences-outlook-2026#:~:92%25...

[39] https://www.slalom.com/us/en/insights/life-sciences-outlook-2026#:~:62%25...

[40] https://sakaradigital.com/blog/ai-center-of-excellence-2m-budget-mid-cap-pharma/#:~:Month...

[41] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:Featu...

[42] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:Roche...

[43] https://nvidianews.nvidia.com/news/nvidia-and-lilly-announce-co-innovation-lab-to-reinvent-drug-discovery-in-the-age-of-ai#:~:Compa...

[44] https://nvidianews.nvidia.com/news/nvidia-and-lilly-announce-co-innovation-lab-to-reinvent-drug-discovery-in-the-age-of-ai#:~:The%2...

[45] https://www.signifytechnology.com/news/machine-learning-engineer-salary-benchmarks-us-market-2025-2026/#:~:Salar...

[46] https://www.kore1.com/ai-engineer-salary-guide/#:~:AI%20...

[47] https://www.benchling.com/biotech-ai-report-2026#:~:Inter...

[48] https://sakaradigital.com/blog/ai-center-of-excellence-2m-budget-mid-cap-pharma/#:~:Mid-c...

[49] https://checkthat.ai/brands/veeva-systems/pricing#:~:Quick...

[50] https://www.vendr.com/marketplace/veeva-systems#:~:Media...

[51] https://ai-scanner.com/platforms/veeva-vault-clinical#:~:Veeva...

[52] https://www.zs.com/insights/scaling-ai-in-pharma-cdio-2026#:~:CIOs%...

[53] https://itbudgetcalculator.com/#:~:Indus...

[54] https://itbudgetcalculator.com/#:~:Of%20...

[55] https://www.union.ai/case-study/rezo-accelerates-drug-discovery-while-saving-90-on-compute-costs-with-union-ai#:~:Massi...

[56] https://www.union.ai/case-study/rezo-accelerates-drug-discovery-while-saving-90-on-compute-costs-with-union-ai#:~:We%20...

[57] https://www.union.ai/case-study/rezo-accelerates-drug-discovery-while-saving-90-on-compute-costs-with-union-ai#:~:to%20...

[58] https://www.nvidia.com/en-eu/case-studies/generative-ai-for-small-molecule-drug-discovery/#:~:Terra...

[59] https://aws.amazon.com/solutions/case-studies/tangram-case-study/#:~:300x%...

[60] https://aws.amazon.com/solutions/case-studies/tangram-case-study/#:~:1%2C0...

[61] https://pharma.economictimes.indiatimes.com/amp/news/pharma-industry/50-of-indian-pharma-companies-exploring-ai-driven-solutions-ey-report/118330020#:~:50%20...

[62] https://pharma.economictimes.indiatimes.com/amp/news/pharma-industry/50-of-indian-pharma-companies-exploring-ai-driven-solutions-ey-report/118330020#:~:Indus...

[63] https://pharma.economictimes.indiatimes.com/amp/news/pharma-industry/50-of-indian-pharma-companies-exploring-ai-driven-solutions-ey-report/118330020#:~:75%20...

[64] https://aws.amazon.com/solutions/case-studies/genentech-generativeai-case-study/#:~:This%...

[65] https://aws.amazon.com/solutions/case-studies/genentech-generativeai-case-study/#:~:total...

[66] https://www.fiercepharma.com/marketing/two-thirds-top-20-pharmas-have-banned-chatgpt-and-many-life-sci-call-ai-overrated-survey#:~:thoug...

[67] https://www.fiercepharma.com/marketing/two-thirds-top-20-pharmas-have-banned-chatgpt-and-many-life-sci-call-ai-overrated-survey#:~:many%...

[68] https://www.fiercepharma.com/marketing/two-thirds-top-20-pharmas-have-banned-chatgpt-and-many-life-sci-call-ai-overrated-survey#:~:more%...

[69] https://www.fiercepharma.com/marketing/two-thirds-top-20-pharmas-have-banned-chatgpt-and-many-life-sci-call-ai-overrated-survey#:~:81%25...

[70] https://www.contextwindows.ai/case-study/innophore-drug-discovery#:~:Cost%...

[71] https://intuitionlabs.ai/articles/ai-adoption-pharma-biotech-benchmarks#:~:altho...

[72] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:ident...

[73] https://www.slalom.com/us/en/insights/life-sciences-outlook-2026#:~:AI%20...

[74] https://www.benchling.com/biotech-ai-report-2026#:~:liter...

[75] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:trust...

[76] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:Nearl...

[77] https://www.zs.com/insights/scaling-ai-in-pharma-cdio-2026#:~:is%20...

[78] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:A%20m...

[79] https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development#:~:1.%20...

[80] https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026#:~:Build...

[81] https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026#:~:AI%20...

[82] https://www.gartner.com/en/newsroom/press-releases/2026-02-03-gartner-forecasts-worldwide-it-spending-to-grow-10-point-8-percent-in-2026-totaling-6-point-15-trillion-dollars#:~:Total...

[83] https://www.gartner.com/en/newsroom/press-releases/2026-02-03-gartner-forecasts-worldwide-it-spending-to-grow-10-point-8-percent-in-2026-totaling-6-point-15-trillion-dollars#:~:Proje...

[84] https://blogs.nvidia.com/blog/state-of-ai-report-2026/#:~:Overa...

[85] https://blogs.nvidia.com/blog/state-of-ai-report-2026/#:~:The%2...

[86] https://blogs.nvidia.com/blog/state-of-ai-report-2026/#:~:Overa...

[87] https://blogs.nvidia.com/blog/state-of-ai-report-2026/#:~:Overa...

[88] https://blogs.nvidia.com/blog/state-of-ai-report-2026/#:~:North...

[89] https://blogs.nvidia.com/blog/state-of-ai-report-2026/#:~:More%...

[90] https://blogs.nvidia.com/blog/state-of-ai-report-2026/#:~:North...

[91] https://www.svb.com/trends-insights/reports/healthcare-investments-and-exits/#:~:Nearl...

[92] https://www.prnewswire.com/news-releases/ai-investment-accounted-for-nearly-half-of-healthcare-investment-in-2025-silicon-valley-bank-releases-17th-healthcare-investments-and-exits-report-302656179.html#:~:Healt...

[93] https://www.prnewswire.com/news-releases/ai-investment-accounted-for-nearly-half-of-healthcare-investment-in-2025-silicon-valley-bank-releases-17th-healthcare-investments-and-exits-report-302656179.html#:~:%2430...

[94] https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/global-tech-report-2026-life-sciences.pdf#:~:97%20...

[95] https://www.benchling.com/biotech-ai-report-2026#:~:27%25...

[96] https://aws.amazon.com/solutions/case-studies/pfizer-PACT-case-study/#:~:16%2C...

[97] https://intuitionlabs.ai/articles/pharma-digital-transformation-ai-strategy#:~:reduc...

[98] https://intuitionlabs.ai/articles/pharma-digital-transformation-ai-strategy#:~:The%2...

[99] https://www.roche.com/media/releases/med-cor-2026-03-16#:~:Featu...

[100] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:These...

[101] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:Manag...

[102] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:Cash%...

[103] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:Recur...

[104] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:The%2...

[105] https://ir.recursion.com/news-releases/news-release-details/recursion-reports-fourth-quarter-and-full-year-2025-financial/#:~:In%20...

[106] https://exa.ai/websets/directory/insilico-funding#:~:Since...

[107] https://insilico.com/news/ohz9ozx0t1-insilico-medicine-announces-2025-annual#:~:bring...

[108] https://www.pharmexec.com/view/eli-lilly-insilico-enter-2-billion-research-licensing-agreement-advance-ai-drug-discovery#:~:Insil...

[109] https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development#:~:CDER%...

[110] https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development#:~:1.%20...

[111] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:One%2...

[112] https://www.slalom.com/us/en/insights/life-sciences-outlook-2026#:~:56%25...

[113] https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2026-life-sciences-executive-outlook.html#:~:Cyber...

[114] https://exa.ai/websets/directory/insilico-funding#:~:the%2...

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by Adrien Laurent, a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.