How to Build an AI Center of Excellence in Biotech (2026)

6/26/2026 • 45 min read

ai center of excellence | biotech ai strategy | pharma ai governance | ai coe framework | hub and spoke model | chief ai officer | life sciences ai | fda ai guidance | ai operating model | pharma digital transformation

Executive Summary

An **AI Center of Excellence (CoE)** is the centralized (or hub-and-spoke) team, governance structure, and shared platform that a biotech or pharmaceutical company builds to turn scattered artificial intelligence (AI) pilots into repeatable, compliant, production capability. As of July 2026, the case for building one is no longer speculative: McKinsey's Global Institute estimated in 2023 that generative AI alone could unlock $60 billion to $110 billion a year in economic value for the pharmaceutical and medical products industries ([1] mckinsey.com), with $18 billion to $30 billion of that concentrated in commercial functions alone ([2] mckinsey.com). Yet capturing that value has proven difficult without organizational infrastructure: an IntuitionLabs analysis of the field found that nearly every large life-science organization is experimenting with AI, but fewer than 5% have fully integrated it into core operations ([3] intuitionlabs.ai).

This report defines the AI CoE concept, walks through the three dominant organizational archetypes used in biotech and pharma (the centralized hub, the decentralized spoke, and the hybrid hub-and-spoke model), and lays out a practical framework for governance, staffing, and enterprise scaling. Consulting firm ZS Associates finds that roughly 60% of pharmaceutical companies start with a centralized hub model and about 40% start with a decentralized spoke model, though a hybrid hub-and-spoke structure, combining centralized strategy and governance with business-unit-led use case development, tends to prove the most durable over time, a pattern echoed in broader enterprise AI operating-model research showing more than 50% of companies adopt a centrally-led structure for generative AI even when their general data and analytics function is decentralized ([4] aiassemblylines.com). Deloitte frames the stakes in blunter terms: companies that operationalize AI at scale through an AI Centre of Excellence are seeing up to 20% EBITDA uplift, while those stuck in what Deloitte calls "Use Case Theatre" see fragmented governance and no realized value ([5] deloitte.com).

Governance is inseparable from structure in a regulated industry. The U.S. Food and Drug Administration's Center for Drug Evaluation and Research (CDER) and Center for Biologics Evaluation and Research (CBER), working with the European Medicines Agency (EMA), published 10 guiding principles for AI in drug development that emphasize human-centric design, a risk-based approach, and multidisciplinary expertise ([6] fda.gov). The FDA also issued draft guidance in January 2025 setting out a "risk-based credibility assessment framework" for AI models used to support regulatory decisions ([7] fda.gov), while the EMA's own reflection paper on AI in the medicinal product lifecycle, finalized in September 2024, requires that data acquisition, model training, and validation be documented "in a detailed and fully traceable manner in line with GxP requirements" (ema.europa.eu). A biotech AI CoE therefore has to do more than most industry-agnostic AI CoE playbooks describe: it must own model risk management, data lineage, and regulatory liaison as core functions, not afterthoughts.

The report also surveys real capital commitments that illustrate scale, discussed fully in the Case Studies section below: Sanofi is investing $294 million to expand its AI Center of Excellence in Toronto, expected to add 50 high-skilled AI, machine learning, and data science jobs by 2028 on top of more than 150 existing roles; Merck & Co. signed a multi-year, up-to-$1 billion partnership with Google Cloud in April 2026 to deploy an agentic AI platform, Gemini Enterprise, across its research and development (R&D), manufacturing, commercial, and corporate functions ([8] merck.com); and Novartis founded an AI innovation lab with Microsoft as its strategic data-science partner as early as October 2019, an early marker of the trend ([9] novartis.com). For organizations without large-cap budgets, a documented mid-cap playbook shows a functioning AI CoE can be built on roughly $2 million a year with four to five full-time employees ([10] sakaradigital.com).

The report closes with implementation guidance: how to sequence a CoE launch over 12 to 18 months, how to avoid the "ivory tower" failure pattern in which a centralized team accumulates budget and talent but ships nothing the business can use ([11] aiadvisorypractice.com), and how governance, quality, and regulatory affairs functions need to be embedded from day one rather than bolted on after a model reaches production. Life-sciences organizations considering an implementation partner for CoE-adjacent technology work, such as building governed dashboards on top of Veeva Vault CRM or standing up validation-ready AI tooling, should weigh advisory firms against this same organizational-design lens rather than treating vendor selection as a separate exercise from CoE design.

Introduction and Background

Biotechnology and pharmaceutical companies have invested in data science and informatics groups for decades, but the current wave of interest in a formal "AI Center of Excellence" reflects something new: an attempt to industrialize AI the way earlier generations of pharma industrialized quality management and IT service delivery. A **Center of Excellence (CoE)**, in the general management sense, is a centralized team or function that concentrates expertise, defines standards, and supports the rest of the organization in adopting a capability. An AI CoE applies that model specifically to artificial intelligence and machine learning (ML), and in pharma it typically sits at the intersection of R&D, manufacturing, commercial operations, and corporate quality and regulatory functions.

The urgency behind this trend is well documented.McKinsey's March 2025 global survey on the state of AI found that companies with at least $500 million in annual revenue are redesigning workflows and elevating governance faster than smaller organizations, and that more than three-quarters of respondents report using AI in at least one business function ([12] mckinsey.com). At the same time, a McKinsey life-sciences compendium notes that nearly eight in ten companies now use generative AI, yet 80% of them report no tangible bottom-line benefit, a gap McKinsey attributes largely to organizational and governance shortfalls rather than model quality (mckinsey.com). This is the pattern practitioners call "pilot purgatory": proof-of-concept projects proliferate, but few reach validated production use inside a GxP-regulated environment.

Historically, pharma's technology CoEs were narrow and siloed: a biostatistics programming group here, a process-control team there, an informatics unit somewhere else. The rise of deep learning and generative AI after roughly 2020 changed the calculus. AI-native biotechs such as Recursion and Insilico Medicine demonstrated that AI-first drug discovery could compress timelines that traditionally took years, and traditional pharma companies responded with dedicated innovation labs. Novartis founded its AI innovation lab with Microsoft as strategic partner in October 2019, explicitly aiming to bring AI "to the desktop of every Novartis associate" and to tackle generative chemistry and image analysis for personalized therapies ([14] novartis.com). A 2026 peer-reviewed study in the *Journal of Technology Transfer* examined 138 innovation labs established by 24 of the world's 40 largest pharmaceutical companies and found that AI-focused infrastructure is highly concentrated among the largest firms, "suggesting scale-dependent conditions for orchestrating data-intensive capabilities" ([15] link.springer.com). The same study found that labs with structurally centralized AI capability show "more internalised and governance-intensive configurations" than labs where AI is merely embedded within a broader innovation portfolio ([16] link.springer.com).

This report is organized to answer both the strategic question, how should an AI CoE in biotech be structured, and the operational question, what does it take to actually run one under FDA and EMA oversight. Every SECONDARY_QUERY implicit in the brief, how to build an AI CoE generally, the pharma-specific CoE pattern, a usable CoE framework, biotech AI strategy, organizational structure, enterprise scaling, and AI governance, is addressed in a dedicated section below. Where useful, the report distinguishes between vendor claims and independently verified figures, and it flags disagreements between sources rather than smoothing them over.

What an AI Center of Excellence Is, and What It Is Not

A Pharma or biotech AI CoE is best defined by its functions rather than by an org chart box. Drawing on the consensus across ZS, McKinsey, Deloitte, and IntuitionLabs research, an AI CoE in this industry typically performs five roles simultaneously.

- **Strategic alignment**: creating and maintaining a corporate AI roadmap tied to business priorities such as accelerating discovery timelines or reducing manufacturing cost of goods, so that AI investment is not simply "AI for its own sake." Microsoft's Cloud Adoption Framework frames the general-enterprise version of this function directly: "An AI CoE consists of an internal team of experts who drive successful and valuable AI outcomes. The AI CoE prevents fragmented or ungoverned AI" adoption ([17] learn.microsoft.com).
- **Standards and best practices**: defining data-handling conventions, model-development templates, validation procedures, and documentation standards so results are reproducible across teams.
- **Shared infrastructure**: providing common cloud services, MLOps (machine learning operations) tooling, data lakes, and compute so that individual business units are not duplicating platform investment. An enterprise AI operating-model analysis describes this as the "foundational infrastructure" tier of an AI portfolio, the data platforms, model infrastructure, security frameworks, and governance tooling that every initiative depends on and that should be funded as a shared cost rather than folded into individual project budgets ([18] aiassemblylines.com).
- **Governance and oversight**: housing (or coordinating with) the mechanism that reviews AI projects for model risk, regulatory compliance, and ethics, described in detail in the next section.
- **Talent hub**: serving as the recruiting, training, and retention focal point for data scientists, ML engineers, and AI-literate domain experts, an especially acute need given documented shortages of AI-skilled professionals in the sector, discussed further in the Roles and Talent section below.

It is important to distinguish an AI CoE from an innovation lab, a data science team, and a Chief AI Officer (CAIO) mandate, three related but distinct organizational elements that are often conflated in casual usage. An innovation lab, as documented in the 138-lab study cited above, is frequently a physical or semi-autonomous unit oriented toward external partnership and experimentation; a CoE is more often the operating backbone that governs and scales what the lab or individual business units produce. A data science team is a staffing pool; a CoE is a mandate that may or may not own that pool directly. A CAIO is an executive role, increasingly common in pharma since roughly 2023 to 2024, that may sponsor a CoE but is not synonymous with one: the CoE is the delivery mechanism, while the CAIO is frequently the accountable executive who decides whether the CoE reports to the Chief Information Officer (CIO), Chief Digital Officer, or Chief Scientific Officer ([19] sakaradigital.com).

A common failure mode, described bluntly by the AI Advisory Practice's 50-page CoE guide, is the "ivory tower": a centralized team that accumulates talent and budget but delivers nothing the business unit can actually use ([11] aiadvisorypractice.com). The same guide identifies five recurring failure patterns: a talent accumulation trap, a platform-before-strategy mistake, business-unit relationship failures that trigger political resistance, over-centralized governance, and metrics that are misaligned with what executive sponsors actually value. This failure mode is the central design problem the rest of this report addresses: how to get the benefits of centralization (standards, reuse, scarce-talent concentration) without the disconnection from business need that centralization tends to produce.

Organizational Archetypes: Hub, Spoke, and Hub-and-Spoke

The single most consequential design decision for a biotech AI CoE is its organizational archetype. ZS Associates, in a widely cited industry analysis, identifies three archetypes that describe nearly every pharma company's starting point.

The hub (centralized) model places all AI work inside a single enterprise-led team. Business units submit requests to the CoE, which designs, builds, and governs the resulting solutions. Roughly 60% of organizations start here, drawn by plentiful funding, deep technical expertise, and unified executive leadership ([20] zs.com). The risk, per ZS, is that the top-down vision "may become disconnected from on-the-ground business- or function-specific objectives," creating bottlenecks and slower responsiveness to local demand.

The spoke (decentralized) model embeds AI capability directly inside each business unit, a drug development division, a manufacturing site, a commercial team, each fielding its own data science resources. About 40% of organizations start with this pattern ([21] zs.com). It produces solutions closely aligned to local priorities and faster initial prototypes, but at the cost of duplicated effort, inconsistent tooling, and difficulty enforcing enterprise-wide standards.

The hub-and-spoke (hybrid) model combines a central CoE that owns strategy, governance, and reusable capabilities, with individual business units ("spokes") that identify, prioritize, and build tailored use cases. ZS describes this as "more durable in the long run" than either pure archetype ([22] zs.com), and it is the model most large pharma organizations converge toward as they mature past initial pilots.

Table 1 below summarizes the trade-offs among the three archetypes, synthesizing the ZS framework ([23] zs.com) with the general enterprise-scaling logic described above.

Archetype	Structure	Advantages	Challenges
Centralized Hub	Single central AI team; business units submit requests	Plentiful funding, deep expertise, unified vision; fast standardization of methods	Can become disconnected from business needs; risk of bottleneck; slower responsiveness to local demand
Decentralized Spoke	AI teams embedded in each business unit	Solutions align closely with local priorities; faster initial development	Duplication of effort; inconsistent tools and data; reliance on unit-level funding
Hub-and-Spoke Hybrid	Central CoE sets strategy and governance (hub); business-unit AI teams (spokes) build use cases	Balances oversight with agility; allows reuse and scale of best practices	Requires clear role definitions; must manage the tension between control and enablement, and avoid redundant governance layers

Reading this table, the interpretive point is not that one archetype is universally correct but that the archetype should match organizational maturity and scale. An enterprise AI operating-model analysis reaches a parallel conclusion outside pharma specifically: organizations running fewer than 10 active AI initiatives "almost always benefit from a centralized or lightly hub-and-spoke model," while the transition point toward a fuller hub-and-spoke structure is typically when the portfolio exceeds 15 to 20 active initiatives across three or more business units ([24] aiassemblylines.com). The AI Advisory Practice guide adds a third named model, "platform-as-a-service," in which a centralized team transitions to providing an internal platform that autonomous business-unit AI teams consume directly, a pattern that tends to appear only after an organization has already run 15 to 20 active projects through a centralized hub and hit its scaling limits ([25] aiadvisorypractice.com).

Getting started with hub-and-spoke, per ZS, involves four sequential steps: developing a strategic roadmap grounded in a maturity assessment of data infrastructure, technology, talent, and culture; establishing an AI operating model with cross-functional collaboration frameworks for risk management and use-case prioritization; building for value realization through tangible, measured outcomes; and deploying a formal use-case prioritization engine with strategic alignment, a structured intake mechanism, and diverse-stakeholder decision-making ([26] zs.com). Companies that fail to plan this transition risk falling behind as AI's underlying computational power continues to compound faster than most organizations can absorb responsibly, a dynamic Deloitte separately frames as a "business model fault line" rather than a routine technology upgrade cycle ([27] deloitte.com).

Roles, Talent, and Team Structure

An AI CoE in biotech requires a blend of technical and domain expertise that goes beyond a typical enterprise data science team, largely because of the regulatory and scientific specificity of the work. Based on the IntuitionLabs organizational-design analysis and the Sakara Digital mid-cap staffing model, the following roles recur across CoE designs regardless of company size.

- **CoE Director or Head of CoE**: the senior executive accountable for AI strategy and business alignment, typically reporting to the CIO, Chief Digital Officer, or Chief Scientific Officer depending on where the CoE's mandate originated ([28] sakaradigital.com).
- **Principal AI Engineer or Architect**: owns technology-stack decisions and production-grade implementation patterns, serving as the technical authority the rest of the CoE and the business units trust.
- **Data Scientists and ML Engineers**: build and validate models, and, per McKinsey, increasingly need to expand "beyond traditional data science roles to include new skills, AI engineering, large language model (LLM) fine-tuning, and business translation" ([29] intuitionlabs.ai).
- **Data Engineers and Architects**: build the pipelines, feature stores, and MLOps infrastructure that make data accessible and models reproducible.
- **AI Product Manager**: translates business problems into AI use cases, owns the use-case portfolio and value tracking, and serves as the bridge between the CoE and the business functions whose problems it is trying to solve.
- **AI Quality and Governance Lead**: owns the governance framework, validation discipline, regulatory engagement, and defensible documentation; Sakara Digital's mid-cap analysis calls this "the most underweighted role" in CoE staffing precisely because pharma's regulatory environment makes this work "non-negotiable" ([30] sakaradigital.com).
- **Domain Experts**: scientists, clinicians, and regulatory affairs specialists who translate business problems into AI use cases and validate model outputs, particularly important at the spokes in a hub-and-spoke design.

Talent scarcity is a persistent constraint. EY-Parthenon's report with Microsoft lists a shortage of AI-skilled professionals and organizational resistance to change among the operational barriers slowing adoption across the life-sciences sector (expresspharma.in). Separate research on the Chief AI Officer role specifically finds that the share of organizations with a CAIO has grown from 11% in 2023 to 26% as of the most recent count, with CAIOs who report directly to the CEO or Chief Operating Officer (COO) associated with measurably better AI-spend outcomes than peers without a dedicated executive owner ([31] aiassemblylines.com). At mid-cap scale, Sakara Digital's model recommends a compact team of 4 to 5 full-time employees, with total staffing cost of $1.2 million out of a $2 million annual CoE budget, at a loaded compensation range of $240,000 to $300,000 per FTE, discussed further in *Table 2* below. By contrast, the AI Advisory Practice's guide describes 12 critical roles for a "fully functioning" large-cap AI CoE, with an explicit sequencing logic for hiring them in phases ([32] aiadvisorypractice.com). The gap between a 5-person mid-cap CoE and a 12-role large-cap CoE illustrates why scope discipline, deciding explicitly what the CoE will and will not do, is the central resourcing decision every biotech has to make before hiring begins.

Governance: Regulatory, Ethical, and Model Risk Frameworks

Governance is the feature that most distinguishes a biotech or pharma AI CoE from its counterpart in, say, retail or financial services. AI systems that touch clinical trial data, drug safety signals, or manufacturing quality control intersect directly with regulatory regimes that were not designed with machine learning in mind, and regulators have moved to close that gap over the past three years.

FDA guidance. The FDA's CDER and CBER, in collaboration with EMA, published 10 guiding principles of good AI practice in drug development, "tailored to the drug development cycle" and intended "to fully realize the potential of AI while ensuring reliability of the information to ensure patient safety and regulatory excellence" ([33] fda.gov). They are: human-centric by 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 ([34] fda.gov), a list the agency situates within "Artificial Intelligence for Drug Development" as a standing CDER program area rather than a one-time guidance release ([35] fda.gov). Separately, the FDA issued draft guidance in January 2025, docketed as FDA-2024-D-4689, titled "Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products," which lays out a risk-based credibility assessment framework tied to a defined "context of use" (COU) for any AI model submitted in support of a regulatory decision ([7] fda.gov). The guidance was issued jointly by an unusually wide set of FDA offices, including the Center for Veterinary Medicine, the Oncology Center of Excellence, the Center for Biologics Evaluation and Research, the Center for Devices and Radiological Health, and the Center for Drug Evaluation and Research ([36] fda.gov), itself a signal that AI oversight in this industry is not confined to a single regulatory center and that a CoE's regulatory-liaison function needs relationships across several FDA centers, not just CDER. Public comments on the guidance remain open on a rolling basis under 21 CFR 10.115(g)(5), meaning the credibility-assessment framework a CoE builds today should be treated as provisional rather than final ([37] fda.gov). This built on an earlier discussion paper ([38] fda.gov), first issued in May 2023 and revised in February 2025, that mapped AI/ML use cases across drug discovery (target identification, compound screening and design), nonclinical research, clinical research (patient recruitment, dose and dosing regimen optimization, adherence, retention, site selection, clinical endpoint assessment), postmarket safety surveillance (case processing, evaluation, and submission), and advanced pharmaceutical manufacturing (process design optimization, advanced process control, smart monitoring and maintenance, trend monitoring) ([39] fda.gov). CDER has publicly noted "a significant increase in the number of drug application submissions using AI components over the past few years," spanning nonclinical, clinical, postmarketing, and manufacturing phases ([40] fda.gov), noting these submissions "traverse the drug product life cycle" ([41] fda.gov), and defines AI broadly as "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments" ([42] fda.gov). CDER further notes machine learning specifically as "a set of techniques that can be used to train AI algorithms to improve performance at a task based on data," the subset of AI most commonly used across the drug product life cycle ([43] fda.gov).

EMA guidance. The EMA's reflection paper on the use of AI in the medicinal product lifecycle, finalized on 9 September 2024 under reference EMA/CHMP/CVMP/83833/2023, spans the full lifecycle from drug discovery through post-authorization surveillance (ema.europa.eu). On manufacturing specifically, the EMA expects AI/ML use in process design, optimization, in-process quality control, and batch release to increase, and directs that model development follow existing quality risk management principles under ICH Q8, Q9, and Q10 "awaiting revision of current regulatory requirements and GMP standards" (ema.europa.eu). On data quality, the paper requires documented consideration of data representativeness and class imbalance before a model is trained. It sets a demanding bar for pivotal clinical trials specifically: "incremental learning approaches are not accepted, and any modification of the model during the trial requires a regulatory interaction to amend the statistical analysis plan," and prior to database lock, the data pipeline and all models must be "frozen and documented in a traceable manner" (ema.europa.eu). The paper is more permissive in lower-risk settings: pharmacovigilance applications, by contrast, "may allow a more flexible approach to AI/ML modelling and deployment, where incremental learning can continuously enhance models" (ema.europa.eu). This differential treatment by risk tier is itself the core design principle a CoE's governance function needs to operationalize: not every model needs the same level of control. Jesper Kjær, Director of the Data Analytics Centre at the Danish Medicines Agency and co-chair of EMA's Big Data Steering Group, framed the underlying regulatory posture when the reflection paper first opened for consultation: "AI brings exciting opportunities to generate new insights and improve processes. To embrace them fully, we will need to be prepared for the regulatory challenges presented by this quickly evolving ecosystem" (ema.europa.eu).

NIST AI Risk Management Framework (AI RMF). Released on 26 January 2023, the voluntary NIST AI RMF is organized around four functions, Govern, Map, Measure, and Manage, intended to help organizations incorporate trustworthiness considerations into the design, development, use, and evaluation of AI systems ([44] nist.gov). The framework was developed through "a consensus-driven, open, transparent, and collaborative process that included a Request for Information, several draft versions for public comments, multiple workshops" ([45] nist.gov), and NIST separately launched a Trustworthy and Responsible AI Resource Center on 30 March 2023 to support implementation and international alignment with the AI RMF ([46] nist.gov). NIST later published a companion Generative AI Profile, NIST-AI-600-1, in July 2024, to help organizations identify risks unique to generative AI ([47] nist.gov). A companion NIST AI RMF Playbook, published alongside the framework, provides implementation-level detail organizations can adapt directly into internal CoE governance procedures ([48] nist.gov). Many life-sciences AI governance programs use the NIST AI RMF as a horizontal baseline and layer FDA- and EMA-specific requirements on top of it, a layering approach IntuitionLabs' own Trust Center describes explicitly: every AI-enabled solution the firm delivers ships with "an intended use statement, data sheet describing training and reference data lineage, a model card, a validation protocol with acceptance criteria, a bias and fairness assessment where relevant, and explicit human-in-the-loop provisions for GxP-impacting decisions," aligned with the NIST AI RMF and FDA AI/ML guidance ([49] intuitionlabs.ai).

Outside the United States and the EMA's own scientific-guideline track, the European Union's AI Act adds a third regulatory layer that a biotech CoE operating in Europe must track directly: USDM Life Sciences' 2026 governance white paper notes that "the EU AI Act's high-risk enforcement provisions take effect August 2026," requiring conformity assessments, technical documentation, and continuous monitoring for qualifying systems ([50] usdm.com). The same white paper flags a governance blind spot specific to validated platform environments: partner platforms are "embedding AI directly into validated GxP environments, often without explicit customer action," which means a CoE's vendor-risk process has to extend to AI features a platform vendor activates inside an already-validated system, not only to AI the CoE builds itself ([51] usdm.com).

Beyond formal regulatory guidance, a peer-reviewed governance framework by Bodnari and colleagues (cited extensively in the IntuitionLabs organizational-design analysis) argues that governance must "require a structured assessment of the underlying problem and a clear justification for why AI is the appropriate tool," to prevent what the authors call "solutionitis," the application of AI without a genuine need ([52] intuitionlabs.ai). Independent peer-reviewed research reaches a similar conclusion from the clinical side: a 2026 systematic review of 35 published healthcare AI governance frameworks in *npj Digital Medicine* found that "inadequately governed AI tools may inadvertently perpetuate systemic inequities, as evidenced by algorithms that unintentionally deprioritize care for underserved populations," and that "effective governance is essential to ensure that technological advancements in healthcare are ethically implemented, efficiently utilized, and consistently prioritize patient welfare" ([53] pmc.ncbi.nlm.nih.gov). The same review notes that while most U.S. hospitals now deploy predictive AI models, "only half assess these systems for bias and two-thirds for accuracy," a gap it calls central to the case for structured governance maturity models ([54] pmc.ncbi.nlm.nih.gov).

In practice, an effective CoE governance program typically operationalizes these principles through a tiered risk classification system that maps use cases to review intensity, a cross-functional AI steering committee with legal, quality, and regulatory representation, and a documented model lifecycle covering project initiation, data audit, model development, validation, deployment, and post-deployment monitoring. This lifecycle approach echoes the EMA's insistence on end-to-end risk management and the FDA's emphasis on life cycle management as one of its 10 guiding principles.

Governance maturity models. Beyond point-in-time frameworks, the 2026 *npj Digital Medicine* systematic review cited above proposes a five-level maturity model, the Healthcare AI Governance Readiness Assessment (HAIRA), spanning "Level 1 (Initial/Ad Hoc) to Level 5 (Leading), with specific benchmarks across all seven governance domains" ([55] pmc.ncbi.nlm.nih.gov). The seven domains the review derived from its analysis of 35 published frameworks are organizational structure, problem formulation, algorithm development and training, model evaluation and validation, external algorithm evaluation and selection, deployment and integration, and monitoring and maintenance ([56] pmc.ncbi.nlm.nih.gov). This kind of tiered maturity model is directly transferable to a biotech AI CoE's self-assessment: a newly formed CoE should expect to sit closer to Level 1 or 2 on most domains, and should treat rapid advancement on the organizational-structure and monitoring domains as a higher priority than chasing Level 5 sophistication on every domain simultaneously. Interviews the same review cites with academic medical centers identified "three existing governance phenotypes for AI evaluation: well-defined, emerging, and interpersonal, each shaped by the available resources and infrastructure at individual institutions" ([57] pmc.ncbi.nlm.nih.gov), a typology that maps closely onto the hub, spoke, and hub-and-spoke archetypes discussed earlier in this report even though it was derived independently from hospital rather than pharma settings.

Enterprise Scaling: From Pilot to Production

Building the CoE's organizational chart and governance framework solves only part of the problem; the harder operational question is how to move dozens or hundreds of individual AI pilots from proof-of-concept into validated, monitored production use. Several sources converge on a similar playbook, even though they use different vocabulary.

McKinsey's 2024 operating-model guidance frames scaling as fundamentally a data-centric exercise: "a technical integration model is only part of what is necessary to generate lasting value from gen AI. Companies must also create gen [erative] AI-ready" organizational structures ([58] hkdca.com). The same guidance notes that a recent McKinsey Global Survey found 65% of companies across sizes, geographies, and industries now use generative AI regularly, "twice as many as last year," underscoring how quickly the baseline of casual adoption has moved even as durable operating-model change lags behind it ([59] hkdca.com), a dynamic the same guidance attributes to companies moving "past the honeymoon phase" of generative AI toward harder operational questions ([60] hkdca.com). ZS's four-step roadmap, described above, sequences a maturity assessment before an operating model, and an operating model before formal use-case prioritization. EY-Parthenon's AI Maturity Framework, developed with Microsoft and unveiled at BioAsia 2025, categorizes organizations into three explicit stages: "foundational" (experimenting without scaling), "innovative" (integrated into select functions but not fully optimized), and "transformational" (enterprise-wide, driving competitive differentiation) (expresspharma.in). EY-Parthenon's National Life Sciences Leader for India, Suresh Subramanian, described the framework's intent as helping organizations move from fragmented AI initiatives to enterprise-wide transformation.

Sakara Digital's mid-cap CoE roadmap gives a concrete 18-month sequence that is representative of how a scaling plan is actually staged in practice: months 1 to 3 focus on foundation building (staffing, governance charter, AI inventory, steering committee); months 4 to 9 deliver the first two to three signature initiatives into production with documented validation evidence; months 10 to 15 focus on capability expansion and pattern transfer to subsequent use-case teams; and months 16 to 18 comprise a formal strategic reset that feeds the year-two budget and priorities. Programs that set milestones beyond their actual capacity produce credibility damage when they miss, while programs that stay within capacity produce credibility momentum when they hit, which is precisely why USDM's governance framework recommends a 90-day pathway rather than an openended one for establishing a baseline: define ownership and decision rights, map intended use and test strategy, protect data lineage and access, and track embedded AI risk across platform updates, in that sequence ([61] usdm.com).

Deloitte's "AI operating model" framing adds a scaling lens organized around a three-pillar structure it declines to fully enumerate publicly but summarizes as aligning AI with business priorities rather than only technology agendas, driving reuse and speed across every function, and embedding governance and compliance "by design" rather than retrofitting it ([62] deloitte.com). Deloitte's broader point, that "AI isn't just a technology shift, it's a business model fault line," and that just as enterprise resource planning (ERP) redefined the back office, AI is redefining the front office ([63] deloitte.com), underscores why scaling AI is treated as a board-level and not purely a CIO-level conversation in the companies that have moved furthest.

Data Analysis and Evidence

The quantitative case for a formal AI CoE structure rests on several independently sourced figures that, taken together, paint a consistent picture: value is real and large, but capture rates remain low without organizational infrastructure.

On **economic opportunity**, the McKinsey Global Institute's 2023 estimate of $60 billion to $110 billion in annual economic value for pharma and medical products from generative AI ([1] mckinsey.com) remains the most widely cited figure in the industry as of mid-2026, and McKinsey's own follow-up analysis attributes $18 billion to $30 billion of that total specifically to commercial functions such as marketing and sales enablement ([64] mckinsey.com). Separately, EY-Parthenon and Microsoft's BioAsia 2025 report projects the AI-in-pharmaceuticals market will reach $16.49 billion by 2034, and the AI-driven medical device market will reach $97.07 billion by 2028 (expresspharma.in). These are different measures (economic value creation versus addressable market size for AI products), and the report treats them as complementary rather than directly comparable.

On **realized value versus experimentation**, the gap is stark. As noted in the introduction, industry analysis finds that fewer than 5% of surveyed companies have realized generative AI as a competitive differentiator yielding consistent financial value. USDM Life Sciences' 2026 AI governance white paper puts a similar gap in starker terms specific to this industry, estimating that 70% to 80% of AI pilots in life sciences "fail to reach production," and that organizations using a unified governance framework instead of fragmented controls see a 35% to 45% reduction in compliance program costs ([65] usdm.com) ([66] usdm.com). McKinsey's own life-sciences compendium separately reports that nearly eight in ten companies use generative AI, yet 80% see no tangible bottom-line benefit (mckinsey.com). EY-Parthenon's India-focused survey found a more optimistic figure in one specific market: 75% of chief experience officers (CXOs) in India's life-sciences sector reported that AI adoption led to cost reductions and improved customer satisfaction (expresspharma.in). The report notes this discrepancy honestly: regional and self-reported survey figures (India CXOs) versus global benchmarking figures (McKinsey's 80% no-benefit finding) are not directly comparable, and the gap may reflect genuine regional variance, different survey methodologies, or a self-selection bias among CXOs willing to make public claims of ROI.

On **operating-model impact**, Deloitte's Africa consulting practice reports that companies operationalizing AI at scale through a Centre of Excellence structure see up to 20% EBITDA uplift, contrasted against organizations stuck in "Use Case Theatre" ([67] deloitte.com). This figure is drawn from Deloitte's own client engagement base rather than an independent academic study, and should be read as a vendor-reported benchmark rather than a peer-reviewed statistic.

On **organizational concentration**, the 2026 Springer *Journal of Technology Transfer* study of 138 pharmaceutical innovation labs found that AI-focused labs are concentrated among the largest firms and are geographically distributed with 25 programmes headquartered in Europe, 25 in the Americas, 20 in Asia, and one each in Africa and Oceania, with 14 programmes operating across multiple continents ([68] link.springer.com). The study's coders achieved 88% inter-rater agreement (Cohen's kappa approximately 0.76) on classifying labs by AI centrality, indicating the classification scheme is methodologically robust ([69] link.springer.com). Of the 40 largest pharmaceutical companies screened for the study, 24 were represented in the final dataset; the remaining 16 companies had no innovation labs the researchers could verify from public disclosure, "which may reflect either the genuine absence of such structures or limitations in public disclosure practices" ([70] link.springer.com).

Table 2 below consolidates the budget and staffing data referenced throughout this report, contrasting a documented mid-cap CoE model with the large-cap range cited by the same source.

Dimension	Mid-Cap Pharma ($500M-$5B revenue)	Large-Cap Pharma
Annual enterprise AI budget	$5M-$20M	$50M-$200M ([71] sakaradigital.com)
Typical CoE annual budget	$2M ([10] sakaradigital.com)	Not publicly standardized; larger, multi-team structures
Core CoE staffing	4-5 FTE ([72] sakaradigital.com)	Up to 12 defined roles ([73] aiadvisorypractice.com)
Example real-world investment	N/A	Sanofi: $294M Toronto CoE expansion, 50 new jobs by 2028 on top of 150+ existing roles ([74] fiercepharma.com); Merck: up to $1B multi-year Google Cloud agentic AI partnership ([75] merck.com)

The interpretive takeaway from this table is that CoE investment scales roughly linearly with company size but the underlying organizational logic, focused staffing, a governance lead, a clear scope, and reusable platform investment, holds at every scale. A CoE that tries to replicate a large-cap design on a mid-cap budget will fail on scope discipline; a CoE that under-invests in governance relative to its regulatory exposure will fail on validation quality regardless of budget size.

A methodological note is worth surfacing here, since it explains part of the gap between vendor optimism and realized value discussed above. McKinsey's own commercial-life-sciences research explicitly debunks the assumption that "generative AI, on its own, will deliver the bulk of the value to be created," arguing instead that the organizational and process changes surrounding a model, not the model itself, determine whether pilots convert into durable value ([76] mckinsey.com), an argument the report frames as offering pharma companies "a once-in-a-century opportunity" only if they build the capability to scale it ([77] mckinsey.com). McKinsey's commercial-adoption survey underpinning much of this analysis drew on responses from more than 100 pharmaceutical, biotech, and medical-device leaders surveyed in December 2023 about their generative AI activity, a modest but methodologically transparent sample size worth noting when weighing the resulting percentages against smaller or less transparent industry surveys ([78] mckinsey.com). McKinsey's broader life-sciences compendium frames the strategic stakes for function leaders directly, noting that most function-level gen AI programs still lack the "growth, strategy, and the next chapter" coherence that separates isolated pilots from an integrated enterprise program (mckinsey.com).

Case Studies and Real-World Examples

Sanofi: Scaling an Established AI Center of Excellence in Toronto

Sanofi's Toronto AI Center of Excellence, first established roughly four years before its most recent expansion, illustrates what a mature CoE looks like once it moves past the initial build phase. In May 2026, Sanofi announced a $294 million investment to scale the hub, supported by an up-to-$5 million conditional grant from the Invest Ontario Fund, expected to create 50 new high-skilled jobs in AI, machine learning, and data science by 2028, on top of more than 150 roles already established across cloud computing, data engineering, software development, bioinformatics, and pharmaceutical data science ([80] fiercepharma.com). Sanofi's chief digital officer, Emmanuel Frenehard, framed the investment in strategic terms: "AI is woven into the fabric of how we discover, develop, produce, launch, and support innovative therapies at Sanofi. Our Toronto AI COE exemplifies how we're harnessing artificial intelligence, accelerating our mission to halve the time from discovery to delivery of innovative treatments" ([81] fiercepharma.com). Beyond the CoE itself, Sanofi has used AI to build "digital twins," or virtual patients, to assist safety and efficacy assessments; data from digital twins helped Sanofi win an FDA pediatric label for its enzyme-replacement therapy Xenpozyme ([82] fiercepharma.com), a concrete example of a CoE-adjacent capability translating into a regulatory outcome.

Merck: An Enterprise-Wide Agentic AI Partnership

Merck & Co.'s April 2026 partnership with Google Cloud, valued at up to $1 billion over multiple years, illustrates a different scaling pattern: rather than building all capability in-house, Merck is deploying Google Cloud's Gemini Enterprise agentic AI platform across R&D, manufacturing, commercial, and corporate functions, with Google Cloud engineers embedded alongside Merck teams ([83] merck.com). Merck's chief information and digital officer, Dave Williams, described the deal as "the next phase of our AI journey, extending our longstanding use of advanced technologies into an intelligent agentic ecosystem that will work alongside our teams" ([84] merck.com). This is best understood as a platform-partnership model that sits alongside, rather than replaces, internal CoE governance: the deal covers deployment across the value chain for a workforce of 75,000 employees worldwide ([85] merck.com), but Merck still needs an internal governance function to validate and monitor any agentic workflow that touches GxP-regulated processes.

Novartis: An Early AI Innovation Lab as CoE Precursor

Novartis's AI innovation lab, founded in October 2019 with Microsoft as its strategic AI and data-science partner, predates the current wave of pharma AI CoEs by several years and illustrates the innovation-lab-to-CoE evolutionary path documented in the Springer technology-transfer study. The lab was designed around two core objectives, "AI Empowerment" (bringing AI tools to every Novartis associate's desktop) and "AI Exploration" (tackling hard computational problems such as generative chemistry and cell and gene therapy optimization) ([86] novartis.com). Novartis CEO Vas Narasimhan framed the effort as central to the company's broader digital strategy: "Pairing our deep knowledge of human biology and medicine with Microsoft's leading expertise in AI could transform the way we discover and develop medicines for the world" ([87] novartis.com). Novartis situated the lab within four broader digital strategic priorities, including scaling 12 "digital lighthouse projects" and becoming, in its words, "the #1 partner in the tech ecosystem" ([88] novartis.com), an early articulation of the strategic-alignment function that later CoE frameworks would formalize.

AstraZeneca: Embedding AI Across the R&D Value Chain Rather Than in a Single Lab

AstraZeneca illustrates a fourth pattern: rather than concentrating AI in one lab or CoE brand name, the company describes data science and AI as “embedded across our R&D to enable our scientists to push the boundaries of science to deliver potentially life-changing medicines” ([89] astrazeneca.com). Jim Weatherall, the company’s data-science leadership voice on the page, frames the goal as “transforming R&D, helping us turn science into medicine more quickly and with a higher probability of success” ([90] astrazeneca.com). In 2021, AstraZeneca selected its first two AI-generated drug targets into its portfolio through a collaboration with BenevolentAI, sourced from internal knowledge graphs that integrate genomic, disease, drug, clinical, and safety data ([91] astrazeneca.com). AstraZeneca reports using AI to help deduce the best molecules to make “across 70 percent of our small molecule chemistry projects” ([92] astrazeneca.com), and its AI-assisted pathology image-analysis systems have the potential to cut tissue-sample analysis time by more than 30% ([93] astrazeneca.com). AstraZeneca’s Centre for Genomics Research is separately working toward the analysis of up to two million genomes by 2026, a scale of infrastructure investment consistent with the Springer study’s finding that dedicated AI infrastructure concentrates among the largest firms.

This distributed-embedding pattern is a genuine alternative to the branded-CoE or branded-lab model, and it illustrates why a hub-and-spoke reading of “AI CoE” should not assume every large pharma names its central function a Center of Excellence: the underlying functions, strategic alignment, shared data infrastructure, and governed use-case delivery, can be organized without a single named organizational unit, provided the enterprise data and AI architecture is genuinely shared across R&D groups rather than duplicated by each one.

GSK and Nvidia: A Vendor-Partnered AI Hub for Genomic Drug Discovery

GlaxoSmithKline (GSK) took a different structural approach again ([94] pharmaceutical-technology.com), building a dedicated, physically located AI hub in London and partnering with chipmaker Nvidia to apply computation to drug and vaccine discovery. The lab is designed to “leverage its genetic and genomic data to advance the designing and development process of new medicines and vaccines,” using biomedical data, AI approaches, and advanced computing platforms to obtain genetic and clinical data “with accuracy and scale” ([95] pharmaceutical-technology.com). The GSK case is a useful counterpoint to AstraZeneca’s embedded model: it shows a named, centrally located AI hub anchored by a single infrastructure vendor partnership, closer to the “hub” archetype ZS describes, rather than AstraZeneca’s cross-R&D embedding pattern, underscoring that even within a small sample of named pharma AI programs, no single organizational archetype dominates. Trade press covering the announcement noted the hub was “newly built” at the time and located in London specifically to draw on the UK’s genomics and computational biology talent base ([96] pharmaceutical-technology.com), framing the partnership’s purpose as applying “AI and computation to drug discovery process” ([97] pharmaceutical-technology.com).

A Global Pharma AI Innovation Hub Consolidating Fragmented Data Science (Hypothetical Example)

To illustrate how a fragmentation problem typically resolves in practice, consider a composite scenario drawn from consulting case patterns documented across the industry: a large biopharmaceutical company faces siloed data science efforts spread independently across R&D, clinical trials, and manufacturing, with no unifying strategy. **(Hypothetical Example)** The company engages advisors to build a centralized AI Innovation Hub that defines a corporate AI roadmap and governance model, standardizes data preparation and model validation processes, accelerates discovery use cases such as predictive analytics for candidate molecules and clinical trial patient-recruitment models, extends AI into manufacturing and supply-chain optimization such as demand forecasting and predictive maintenance, aligns governance with regulatory requirements including validation protocols and security measures, and rolls out AI-literacy training across the organization to build a culture of adoption ([98] intuitionlabs.ai). The expected outcomes in this pattern, faster drug discovery, streamlined trials, and improved operational efficiency, mirror the general logic that recurs across every real-world case examined in this report: strategy and governance first, use-case acceleration second, and culture-building throughout.

Implications and Future Directions

Several structural trends will shape how biotech AI CoEs evolve over the next 12 to 24 months. First, the shift from predictive and generative AI toward **agentic AI**, systems that can autonomously plan and execute multi-step workflows, is already underway at scale, exemplified by Merck's Gemini Enterprise deployment and Roche's reported deployment of additional Nvidia Blackwell graphics processing units (GPUs) to expand what it describes as the industry's largest "AI factory" ([99] fiercepharma.com). Agentic systems raise the governance stakes considerably: the EMA's principle that any model modification during a pivotal trial "requires a regulatory interaction" was written with relatively static models in mind, and CoEs will need updated governance mechanisms for agents that adapt behavior across a workflow rather than producing a single static prediction.

Second, the **Chief AI Officer** role is becoming a more common anchor point for CoE sponsorship, though its authority and reporting line remain unsettled across the industry. As documented in the Sakara Digital CAIO playbook, the role's mandate varies by company, some CAIOs are hired to drive AI-enabled drug discovery, others to lead manufacturing and supply-chain AI, others to coordinate enterprise-wide strategy and governance, and the first 90 days in the role typically determine whether it becomes "a transformative leadership position or a high-profile coordination function with limited authority" ([100] sakaradigital.com). The same playbook sequences the role's early work into distinct phases, mandate clarification in days 1 to 15, an AI inventory and stakeholder mapping exercise in days 15 to 30, governance setup anchored in days 30 to 60, and signature-initiative selection in days 45 to 75, before board-level positioning closes out days 60 to 90 ([101] sakaradigital.com), a sequence that mirrors the CoE-build roadmap described in the Enterprise Scaling section above and reinforces that the CAIO's first task is usually inventory and governance before any new AI capability is commissioned. The playbook closes with a discussion of "the pitfalls that derail CAIOs and how to avoid them," a reminder that organizational authority, not technical skill, is the more common failure point for the role ([102] sakaradigital.com).

Third, regulatory frameworks will continue to mature. The FDA's January 2025 draft guidance remains in draft status as of this report's publication and is expected to move toward finalization, while a 2026 peer-reviewed critical review in the *Journal of Chemistry* has already scrutinized the draft's risk-based credibility assessment framework for gaps and ambiguities ([103] onlinelibrary.wiley.com). CoEs that build governance processes anchored too rigidly to a single draft guidance risk needing significant rework once final guidance issues; the more durable approach is to build governance around the underlying principles (risk-based tiering, traceable data lineage, human oversight, life cycle management) that are unlikely to change even if procedural specifics do.

A related governance-readiness gap shows up in broader enterprise AI research: one 2026 industry survey found governance readiness sitting at only 30% among companies already deploying AI in production, compared with 43% for technical infrastructure readiness and 40% for data-management readiness, which an analysis of the finding calls "not a technology gap. It's an operating model gap" ([104] aiassemblylines.com). This pattern, technical and data infrastructure maturing faster than the governance function meant to oversee it, is precisely the imbalance a well-designed AI CoE governance lead is meant to correct before an organization scales its AI footprint further.

Fourth, **industry concentration** in AI capability is likely to persist. The Springer technology-transfer study's finding that AI-focused innovation infrastructure is concentrated among the largest pharmaceutical firms, and its explicit “scale-dependent” framing, suggests that mid-cap and smaller biotechs will increasingly need external partnerships, whether through consulting advisors, cloud hyperscaler platforms, or specialized vendors, to access CoE-equivalent capability without large-cap-scale budgets. This is consistent with the mid-cap $2 million CoE model’s heavy reliance on commercial software-as-a-service (SaaS) tooling and external partnership budget rather than custom-built infrastructure, as described in the Roles and Talent section above.

For life-sciences organizations evaluating how to close capability gaps around CoE-adjacent technology work, the same organizational-design discipline applies whether the work is building AI governance packets, standing up validated analytics on top of existing systems such as Veeva Vault CRM, or advising on quality-system integration; advisory relationships should be assessed against the same risk-tiering and validation rigor a CoE would apply to any internal build, rather than treated as a procurement decision separate from AI governance strategy.

IntuitionLabs’ own Trust Center frames the underlying philosophy compactly: “In pharmaceutical and biotech, a vendor becomes part of your validated system boundary,” which is why the firm builds IQ/OQ/PQ (installation, operational, and performance qualification) and audit evidence into engagements “from day one, not bolted on after delivery” ([105] intuitionlabs.ai) ([106] intuitionlabs.ai). That same logic, that AI governance obligations travel with whichever system boundary the AI sits inside, applies whether the system is an internally built CoE model or an externally sourced platform feature.

Frequently Asked Questions (FAQs)

What is an AI Center of Excellence in biotech?
It is a centralized team, governance framework, and set of shared platforms that a biotech or pharmaceutical company establishes to coordinate AI strategy, standards, infrastructure, and oversight across business units, moving AI from isolated pilots to production-scale, regulatorily defensible deployment, as defined in the introduction above.

How do you build an AI center of excellence generally, outside biotech?
The general pattern, per Microsoft’s Cloud Adoption Framework, is that “an AI CoE consists of an internal team of experts who drive successful and valuable AI outcomes,” preventing fragmented or ungoverned AI efforts by centralizing standards while enabling distributed execution ([17] learn.microsoft.com). In biotech, this general pattern is overlaid with GxP-specific validation, data lineage, and regulatory-liaison requirements that a horizontal enterprise CoE would not need.

What is a pharma AI center of excellence, specifically?
Microsoft’s Cloud Adoption Framework, which describes how to build an AI CoE “in your organization” generally ([107] learn.microsoft.com), states that a CoE establishes “a strong foundation for AI initiatives” and provides “business and technical consultation that supports successful AI integration” ([108] learn.microsoft.com); the pharma-specific version is the same general concept applied to an organization operating under FDA, EMA, and equivalent regulatory oversight, meaning the CoE’s governance function must integrate directly with existing quality management systems (QMS) under GxP standards, treating AI models and code as controlled documents subject to version control and change management, similar to how a laboratory analytical method would be validated ([109] intuitionlabs.ai).

Is there a standard AI CoE framework?
No single universally adopted framework exists, but there is broad convergence across McKinsey, Deloitte, ZS, and independent guides around three structural choices (centralized hub, decentralized spoke, hub-and-spoke hybrid), a common set of governance functions (data governance, model validation, ethics and fairness, accountability and oversight), and a staged maturity progression from foundational experimentation to enterprise-wide transformation (expresspharma.in).

What should a biotech AI strategy include beyond the CoE itself?
A defensible strategy distinguishes between strategy (where AI will create durable competitive advantage) and roadmap (the executable implementation plan), a distinction the source frames sharply: "most of what gets produced is closer to a wish list than a strategy" ([110] sakaradigital.com); one industry tracking estimate suggests roughly 70% of pharma AI strategies fail to produce a corresponding executable roadmap within 12 months of board approval, which the source calls the largest single source of AI investment underperformance in the industry ([111] sakaradigital.com). Readers should treat this as a single consultancy's internal tracking figure rather than an independently audited industry statistic.

How does AI governance in biotech differ from generic corporate AI governance?
Generic AI governance programs, including NIST's AI RMF, are horizontally applicable across industries. Biotech-specific governance layers on top of that baseline: FDA's context-of-use and risk-based credibility assessment requirements ([112] fda.gov), EMA's GxP-aligned data traceability requirements (ema.europa.eu), and patient-safety-specific fairness and bias mitigation requirements for any model influencing clinical decisions.

How much does a biotech AI CoE cost to run?
Costs scale with company size: a documented mid-cap model runs approximately $2 million a year for a 4-to-5-person team, while large-cap enterprise AI budgets, of which the CoE is one part, range from $50 million to $200 million annually, per *Table 2* above. Individual capital commitments at the largest companies, such as Sanofi's $294 million Toronto expansion or Merck's up-to-$1-billion Google Cloud partnership, illustrate the upper bound of what enterprise-scale AI investment can look like.

Conclusion

Building an AI Center of Excellence in biotech is fundamentally an organizational design problem wrapped inside a regulatory compliance problem. The technology itself, generative models, machine learning pipelines, and increasingly agentic systems, is broadly similar to what other regulated and unregulated industries use. What makes the biotech case distinctive is that governance cannot be treated as a downstream compliance checkbox: FDA's 10 guiding principles, EMA's GxP-traceable data lifecycle requirements, and NIST's four-function risk management framework all point toward the same underlying discipline, embed risk classification, human oversight, and documentation into the AI development process from day one rather than retrofitting it once a model reaches production.

The organizational evidence converges on a similar conclusion from a different angle. Roughly 60% of companies start with a centralized hub and 40% with a decentralized spoke, but the pattern that proves durable over time is a hybrid hub-and-spoke model that concentrates strategy, governance, and reusable platforms centrally while letting business units own use-case execution close to the domain expertise that AI needs to be useful. This is true whether the company in question is a mid-cap biotech running a lean $2 million CoE with five people, or a top-10 pharma company like Sanofi or Merck deploying hundreds of millions to billions of dollars into enterprise-wide AI infrastructure.

The gap between AI experimentation and realized value, with McKinsey data suggesting fewer than 5% of companies have moved AI from pilot to durable competitive differentiator, is not primarily a model-quality problem. It is an organizational-design and governance problem, and it is precisely the problem an AI Center of Excellence, correctly scoped, staffed, and governed, is built to solve. Organizations beginning this journey in the second half of 2026 have more documented precedent to draw on than their predecessors did even two years earlier: named case studies, published regulatory guidance, peer-reviewed organizational research, and detailed budget models across company sizes now exist where only ad hoc consulting frameworks existed before. The remaining execution risk lies not in a lack of available guidance but in the discipline to scope a CoE honestly against an organization's actual maturity, budget, and regulatory exposure, rather than aspiring to a large-cap design the organization cannot sustain.

External Sources

[1] https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality#:~:The%2...

[2] https://www.mckinsey.com/industries/life-sciences/our-insights/early-adoption-of-generative-ai-in-commercial-life-sciences#:~:The%2...

[3] https://intuitionlabs.ai/pdfs/pharma-ai-center-of-excellence-org-design-scaling.pdf#:~:nearl...

[4] https://aiassemblylines.com/resources/ai-initiatives-operating-model#:~:McKin...

[5] https://www.deloitte.com/za/en/services/consulting/perspectives/the-rise-of-the-ai-operating-model.html#:~:compa...

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

[7] https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological#:~:this%...

[8] https://www.merck.com/news/merck-and-google-cloud-partner-to-accelerate-agentic-ai-enterprise-transformation/#:~:This%...

[9] https://www.novartis.com/news/media-releases/novartis-and-microsoft-announce-collaboration-transform-medicine-artificial-intelligence#:~:Novar...

[10] https://sakaradigital.com/blog/ai-center-of-excellence-2m-budget-mid-cap-pharma/#:~:%242M...

[11] https://aiadvisorypractice.com/white-papers/ai-coe-guide#:~:Most%...

[12] https://www.hkdca.com/wp-content/uploads/2024/09/data-leaders-operating-guide-scaling-ai-mckinsey.pdf#:~:To%20...

[13] https://www.mckinsey.com/\~/media/mckinsey/industries/life%20sciences/the%20synthesis%20blog/growth%20strategy%20and%20the%20next%20chapter%20in%20life%20sciences/mck250636%20life%20sciences%20compendium_full%20issue-single_v10.pdf#:~:Recen...

[14] https://www.novartis.com/news/media-releases/novartis-and-microsoft-announce-collaboration-transform-medicine-artificial-intelligence#:~:The%2...

[15] https://link.springer.com/article/10.1007/s10961-026-10347-6#:~:AI-fo...

[16] https://link.springer.com/article/10.1007/s10961-026-10347-6#:~:Labs%...

[17] https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/ai/center-of-excellence#:~:An%20...

[18] https://aiassemblylines.com/resources/ai-initiatives-operating-model#:~:The%2...

[19] https://sakaradigital.com/blog/first-90-days-pharma-chief-ai-officer-playbook/#:~:Diffe...

[20] https://www.zs.com/insights/organizing-scaled-ai-pharma-hub-and-spoke-ai-model#:~:AI%20...

[21] https://www.zs.com/insights/organizing-scaled-ai-pharma-hub-and-spoke-ai-model#:~:AI%20...

[22] https://www.zs.com/insights/organizing-scaled-ai-pharma-hub-and-spoke-ai-model#:~:While...

[23] https://www.zs.com/insights/organizing-scaled-ai-pharma-hub-and-spoke-ai-model#:~:AI%20...

[24] https://aiassemblylines.com/resources/ai-initiatives-operating-model#:~:Organ...

[25] https://aiadvisorypractice.com/white-papers/ai-coe-guide#:~:disad...

[26] https://www.zs.com/insights/organizing-scaled-ai-pharma-hub-and-spoke-ai-model#:~:Step%...

[27] https://www.deloitte.com/za/en/services/consulting/perspectives/the-rise-of-the-ai-operating-model.html#:~:AI%20...

[28] https://sakaradigital.com/blog/first-90-days-pharma-chief-ai-officer-playbook/#:~:Diffe...

[29] https://intuitionlabs.ai/pdfs/pharma-ai-center-of-excellence-org-design-scaling.pdf#:~:teams...

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

[31] https://aiassemblylines.com/resources/ai-initiatives-operating-model#:~:IBM%2...

[32] https://aiadvisorypractice.com/white-papers/ai-coe-guide#:~:the%2...

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

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

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

[36] https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological#:~:Guida...

[37] https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological#:~:You%2...

[38] https://www.fda.gov/media/167973/download#:~:Discu...

[39] https://www.fda.gov/media/167973/download#:~:II.%2...

[40] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:CDER%...

[41] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:These...

[42] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:Artif...

[43] https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development#:~:ML%20...

[44] https://www.nist.gov/itl/ai-risk-management-framework#:~:NIST%...

[45] https://www.nist.gov/itl/ai-risk-management-framework#:~:Relea...

[46] https://www.nist.gov/itl/ai-risk-management-framework#:~:On%20...

[47] https://www.nist.gov/itl/ai-risk-management-framework#:~:On%20...

[48] https://www.nist.gov/itl/ai-risk-management-framework#:~:A%20c...

[49] https://intuitionlabs.ai/trust#:~:For%2...

[50] https://www.usdm.com/resources/white-papers/ai-governance-for-life-sciences-enterprise-framework#:~:The%2...

[51] https://www.usdm.com/resources/white-papers/ai-governance-for-life-sciences-enterprise-framework#:~:Partn...

[52] https://intuitionlabs.ai/pdfs/pharma-ai-center-of-excellence-org-design-scaling.pdf#:~:gover...

[53] https://pmc.ncbi.nlm.nih.gov/articles/PMC13004926/#:~:inade...

[54] https://pmc.ncbi.nlm.nih.gov/articles/PMC13004926/#:~:only%...

[55] https://pmc.ncbi.nlm.nih.gov/articles/PMC13004926/#:~:HAIRA...

[56] https://pmc.ncbi.nlm.nih.gov/articles/PMC13004926/#:~:organ...

[57] https://pmc.ncbi.nlm.nih.gov/articles/PMC13004926/#:~:Inter...

[58] https://www.hkdca.com/wp-content/uploads/2024/09/data-leaders-operating-guide-scaling-ai-mckinsey.pdf#:~:Yet%2...

[59] https://www.hkdca.com/wp-content/uploads/2024/09/data-leaders-operating-guide-scaling-ai-mckinsey.pdf#:~:A%20r...

[60] https://www.hkdca.com/wp-content/uploads/2024/09/data-leaders-operating-guide-scaling-ai-mckinsey.pdf#:~:compa...

[61] https://www.usdm.com/resources/white-papers/ai-governance-for-life-sciences-enterprise-framework#:~:Gover...

[62] https://www.deloitte.com/za/en/services/consulting/perspectives/the-rise-of-the-ai-operating-model.html#:~:Align...

[63] https://www.deloitte.com/za/en/services/consulting/perspectives/the-rise-of-the-ai-operating-model.html#:~:Just%...

[64] https://www.mckinsey.com/industries/life-sciences/our-insights/early-adoption-of-generative-ai-in-commercial-life-sciences#:~:with%...

[65] https://www.usdm.com/resources/white-papers/ai-governance-for-life-sciences-enterprise-framework#:~:70%E2...

[66] https://www.usdm.com/resources/white-papers/ai-governance-for-life-sciences-enterprise-framework#:~:35%E2...

[67] https://www.deloitte.com/za/en/services/consulting/perspectives/the-rise-of-the-ai-operating-model.html#:~:compa...

[68] https://link.springer.com/article/10.1007/s10961-026-10347-6#:~:Geogr...

[69] https://link.springer.com/article/10.1007/s10961-026-10347-6#:~:Inter...

[70] https://link.springer.com/article/10.1007/s10961-026-10347-6#:~:This%...

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

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

[73] https://aiadvisorypractice.com/white-papers/ai-coe-guide#:~:the%2...

[74] https://www.fiercepharma.com/pharma/sanofi-expands-ai-capabilities-investing-294m-scale-toronto-hub#:~:is%20...

[75] https://www.merck.com/news/merck-and-google-cloud-partner-to-accelerate-agentic-ai-enterprise-transformation/#:~:This%...

[76] https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality#:~:Gener...

[77] https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality#:~:The%2...

[78] https://www.mckinsey.com/industries/life-sciences/our-insights/early-adoption-of-generative-ai-in-commercial-life-sciences#:~:At%20...

[79] https://www.mckinsey.com/\~/media/mckinsey/industries/life%20sciences/the%20synthesis%20blog/growth%20strategy%20and%20the%20next%20chapter%20in%20life%20sciences/mck250636%20life%20sciences%20compendium_full%20issue-single_v10.pdf#:~:Growt...

[80] https://www.fiercepharma.com/pharma/sanofi-expands-ai-capabilities-investing-294m-scale-toronto-hub#:~:This%...

[81] https://www.fiercepharma.com/pharma/sanofi-expands-ai-capabilities-investing-294m-scale-toronto-hub#:~:%E2%8...

[82] https://www.fiercepharma.com/pharma/sanofi-expands-ai-capabilities-investing-294m-scale-toronto-hub#:~:Data%...

[83] https://www.merck.com/news/merck-and-google-cloud-partner-to-accelerate-agentic-ai-enterprise-transformation/#:~:will%...

[84] https://www.merck.com/news/merck-and-google-cloud-partner-to-accelerate-agentic-ai-enterprise-transformation/#:~:%E2%8...

[85] https://www.merck.com/news/merck-and-google-cloud-partner-to-accelerate-agentic-ai-enterprise-transformation/#:~:help%...

[86] https://www.novartis.com/news/media-releases/novartis-and-microsoft-announce-collaboration-transform-medicine-artificial-intelligence#:~:1%29%...

[87] https://www.novartis.com/news/media-releases/novartis-and-microsoft-announce-collaboration-transform-medicine-artificial-intelligence#:~:Pairi...

[88] https://www.novartis.com/news/media-releases/novartis-and-microsoft-announce-collaboration-transform-medicine-artificial-intelligence#:~:3%29%...

[89] https://www.astrazeneca.com/r-d/data-science-and-ai.html#:~:Data%...

[90] https://www.astrazeneca.com/r-d/data-science-and-ai.html#:~:Data%...

[91] https://www.astrazeneca.com/r-d/data-science-and-ai.html#:~:In%20...

[92] https://www.astrazeneca.com/r-d/data-science-and-ai.html#:~:Today...

[93] https://www.astrazeneca.com/r-d/data-science-and-ai.html#:~:This%...

[94] https://www.pharmaceutical-technology.com/news/nvidia-gsk-ai/#:~:US-ba...

[95] https://www.pharmaceutical-technology.com/news/nvidia-gsk-ai/#:~:The%2...

[96] https://www.pharmaceutical-technology.com/news/nvidia-gsk-ai/#:~:Locat...

[97] https://www.pharmaceutical-technology.com/news/nvidia-gsk-ai/#:~:The%2...

[98] https://intuitionlabs.ai/pdfs/pharma-ai-center-of-excellence-org-design-scaling.pdf#:~:Centr...

[99] https://www.fiercepharma.com/pharma/sanofi-expands-ai-capabilities-investing-294m-scale-toronto-hub#:~:Swiss...

[100] https://sakaradigital.com/blog/first-90-days-pharma-chief-ai-officer-playbook/#:~:The%2...

[101] https://sakaradigital.com/blog/first-90-days-pharma-chief-ai-officer-playbook/#:~:We%20...

[102] https://sakaradigital.com/blog/first-90-days-pharma-chief-ai-officer-playbook/#:~:We%20...

[103] https://onlinelibrary.wiley.com/doi/10.1155/joch/5202999#:~:This%...

[104] https://aiassemblylines.com/resources/ai-initiatives-operating-model#:~:That%...

[105] https://intuitionlabs.ai/trust#:~:In%20...

[106] https://intuitionlabs.ai/trust#:~:We%20...

[107] https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/ai/center-of-excellence#:~:This%...

[108] https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/ai/center-of-excellence#:~:It%20...

[109] https://intuitionlabs.ai/pdfs/pharma-ai-center-of-excellence-org-design-scaling.pdf#:~:A%20l...

[110] https://sakaradigital.com/blog/ai-strategy-roadmap-pharma-executives/#:~:Most%...

[111] https://sakaradigital.com/blog/ai-strategy-roadmap-pharma-executives/#:~:%7E70...

[112] https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological#:~:this%...

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at https://intuitionlabs.ai/contact for a consultation.

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.