IntuitionLabs

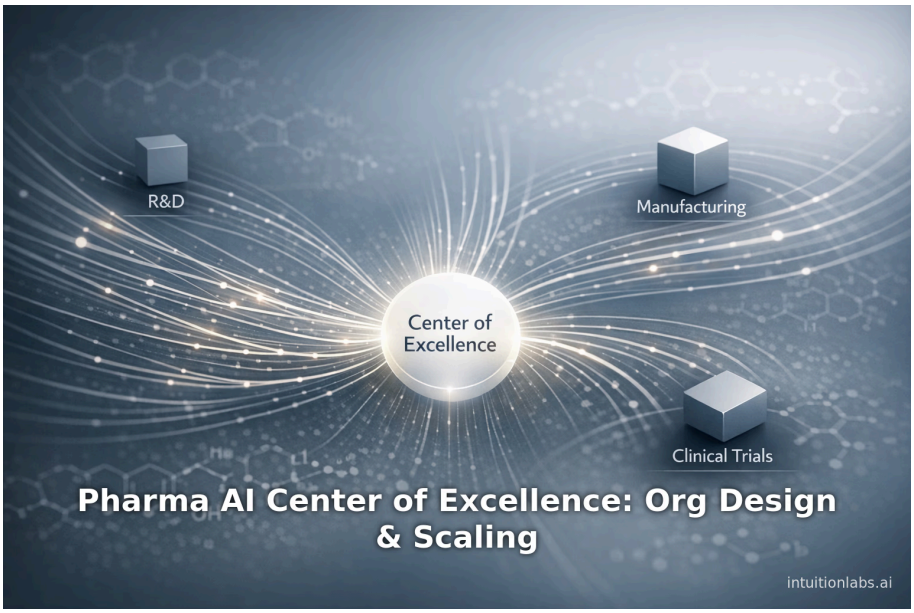
INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Pharma AI Center of Excellence: Org Design & Scaling

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

Published by IntuitionLabs · 2026-04-10 · pharma ai · center of excellence · ai governance · organizational design · enterprise scaling · hub and spoke model · ai operating model · digital transformation

Original article: https://intuitionlabs.ai/articles/pharma-ai-center-excellence-org-design-scaling



In this article

1. Executive Summary
2. Introduction and Background
3. 1. Defining a Pharma AI Center of Excellence
4. Historical Context and Evolution in Pharma
5. 2. Organizational Design
6. 2.1 Organizational Models: Hub, Spoke, and Hybrid
7. 2.2 Roles, Teams, and Talent
8. 2.3 Ways of Working
9. 3. Governance: Managing Risk and Ethics
10. 3.1 Core Principles of AI Governance
11. 3.2 Regulatory Considerations
12. 3.3 Frameworks and Lifecycles
13. 4. Enterprise Scaling Playbook
14. 4.1 Key Strategies for Scaling AI
15. 4.2 Phased Rollout and Pilots
16. 4.3 Metrics and Value Tracking
17. 5. Case Studies and Examples
18. 6. Implications and Future Directions
19. Conclusion

Executive Summary

The pharmaceutical industry stands at a pivotal inflection point as artificial intelligence (AI) promises to reshape every aspect of drug discovery, development, manufacturing, and commercialization. Yet turning AI's potential into enterprise-scale impact requires more than isolated pilots: it demands coordinated strategy, specialized talent, robust governance, and investment in infrastructure. Industry leaders are increasingly converging on the concept of an AI **Center of Excellence (CoE)** – a centralized entity that orchestrates AI capability across the organization. A well-designed Pharma AI CoE can align AI initiatives with business strategy, standardize best practices, and provide oversight for data quality and regulatory compliance. Critical to success are the *organizational design* (how the CoE is structured and staffed), *governance* (the risk, ethics, and regulatory oversight of AI development and use), and an *enterprise-scaling playbook* (a step-by-step process for moving AI from pilot to production across the value chain).

This report provides an in-depth examination of these facets, drawing on academic research, industry analyses, and case examples. It reviews AI organizational archetypes in pharma (centralized "hub", decentralized "spoke", and hybrid models) and their pros and cons (www.zs.com). It outlines the key roles and capabilities an AI CoE needs (data scientists, AI engineers, domain experts, etc.) (www.mckinsey.com). It examines governance requirements unique to healthcare and life sciences, including integration with existing quality systems, data privacy, model validation, and ethical guidelines (www.nature.com) (www.mckinsey.com). It then lays out strategies for scaling AI across an enterprise, based on recent studies and frameworks

(www.mckinsey.com). Throughout, the report cites current examples: from a consulting case of a global pharma building an AI CoE (peymancx.com) to industry reports on AI adoption challenges and maturity levels (www.mckinsey.com) (www.ey.com). The conclusion discusses future trends – generative AI, evolving regulations, and talent needs – that will shape Pharma AI CoEs going forward.

Introduction and Background

The pharmaceutical sector has begun to invest heavily in AI, driven by the promise of speeding drug discovery, optimizing clinical trials, and streamlining operations. A **McKinsey Global Institute** study estimates that generative AI alone could unlock *$60–$110 billion per year* in economic value for the pharmaceutical and medical products industries (www.mckinsey.com). Advances in machine learning (ML) are poised to transform tasks such as target identification, chemical property prediction, patient recruitment, and regulatory filing. In strategic surveys, virtually all large life science organizations report experimenting with AI, and nearly one-third have begun scaling it beyond pilots (www.mckinsey.com).

Yet despite this enthusiasm, moving from experimentation to *enterprise-scale* impact remains challenging. McKinsey finds that only about **5%** of surveyed companies had “realized gen AI as a competitive differentiator” yielding consistent financial value (www.mckinsey.com). The common story is “pilot purgatory”: many proof-of-concept projects show promise, but few reach production. This gap stems not from lack of technical potential, but from organizational, cultural, and data-related barriers. Fragmented AI strategies (often ad-hoc and siloed) abound, and companies frequently lack the infrastructure, talent mix, and governance processes to deploy and maintain AI at scale (www.mckinsey.com) (www.ey.com). Regulatory ambiguity and risk aversion further slow adoption in this highly regulated field (www.ey.com) (www.pharmexec.com).

To overcome these hurdles, leading thinkers emphasize the need for a **structured, enterprise-level approach**. In this context, many pharma firms are establishing or planning an AI Center of Excellence. A CoE acts as a centralized hub (often sponsored by C-level executives) that sets strategy, develops standards, and supports business units in deploying AI solutions. It balances the competing demands of *innovation versus control*: enabling agile use-case development while ensuring consistency, reuse, and compliance. In its own case study, Peyman describes creating a global AI Centre of Excellence for an unnamed biopharmaceutical client, with a structured approach to AI adoption (Peyman).

Notably, analysts at **ZS Associates** describe that most pharma companies begin their AI journeys with one of two archetypes: a centralized “hub” model or a decentralized “spoke” model (www.zs.com). However, a **hybrid “hub-and-spoke”** model – combining a strategic central CoE with empowered business-unit teams – is emerging as the more durable long-term solution (www.zs.com). This report will examine these organizational models in depth. It will then explore how governance (covering both technical and ethical risk management) must be tailored for pharma. Finally, it will outline a playbook for scaling AI from niche projects to enterprise-wide transformation. Throughout, the focus is on the specific needs and context of the pharmaceutical industry – characterized by long R&D cycles, stringent quality requirements, and global regulatory oversight – while drawing on broader AI/innovation literature. McKinsey’s late-summer-2024 survey of more than 100 pharma and medtech leaders found that 5% reported gen AI as a competitive differentiator generating consistent, significant financial value; 32% reported taking steps to scale the technology (McKinsey).

1. Defining a Pharma AI Center of Excellence

A **Center of Excellence (CoE)** is not a new concept; organizations have long used CoEs (for example in quality management, IT, or data analytics) to centralize expertise and set standards. An AI CoE specifically serves to *"operationalize artificial intelligence at scale"* across the enterprise (premainstream.com) (hbr.org). In pharma, an AI CoE has several purposes:

- **Strategic Alignment:** It creates and maintains a corporate AI strategy and roadmap, ensuring AI initiatives align with business priorities (e.g. accelerating discovery or reducing manufacturing costs) (peymancx.com) (www.ey.com). For instance, an EY/Microsoft study proposes an "AI Maturity Framework" to help firms move from fragmented pilots to enterprise-wide transformation (www.ey.com).
- **Standards and Best Practices:** The CoE defines best practices for data handling, model development, validation, and deployment. This might include standardized coding templates, ML pipelines, and documentation standards to ensure reproducibility and quality.
- **Shared Infrastructure:** It can provide shared platforms (cloud services, MLOps tools, data lakes, compute resources) that all teams can use. By centralizing technology investments, reuse is maximized and "scale" becomes practical (www.mckinsey.com).
- **Governance and Oversight:** The CoE often houses the governance function for AI (see Section 3). This means overseeing model risk, compliance with regulations, and ethical guidelines.
- **Talent Hub:** It becomes the focal point for AI skills recruitment, training, and retention. The CoE may staff data scientists, ML engineers, and domain specialists who can be staffed across projects.

In short, a Pharma AI CoE serves as both the **brain and backbone** of the AI transformation. It ensures investment is sustained, duplicative efforts are reduced, and knowledge is shared. Established players in other industries confirm that without such a centralized approach, AI efforts often remain disconnected or fail to add lasting value (www.mckinsey.com) (www.ey.com).

Historical Context and Evolution in Pharma

Pharma has a history of technology CoEs (e.g., for biostatistics or digital health). Early efforts to digitize R&D (in the 2000s and 2010s) often birthed data science groups or informatics centers. However, many of these were siloed – for example, a statistical programming group within clinical trials, or a process control group in manufacturing – rather than addressing AI holistically. The rise of AI (specifically deep learning and generative models) around 2020 shifted the urgency. Companies like Novartis, Roche and GSK began publicly announcing AI labs or partnerships (www.novartis.com) (www.mckinsey.com), recognizing AI's penetration was broad (from image analysis to drug formulation).

Simultaneously, digital-native biotechs (e.g. Recursion, Insilico) were founded with AI at their core, demonstrating a model of *AI-first drug R&D*. Traditional pharma executives took note: if new competitors can iterate faster on target discovery and preclinical tests using ML, incumbents must adapt. The pharmaceutical COVID-19 response (rapid vaccine development) further underscored the need for agility.

In McKinsey's late-summer-2024 survey of pharma and medtech leaders, more than two-thirds of respondents said they planned to significantly increase investment in gen AI (McKinsey). However, McKinsey's January 2025 report on US workplaces across industries found that most organizations investing in AI were not getting the returns they had hoped for (McKinsey). Industry opinion leaders began publishing guidelines: e.g. HBR's *"How to Set Up an AI Center of Excellence"* (2019) and later LinkedIn posts on corporate AI CoEs (hbr.org) (www.zartis.com). Professional services (McKinsey, ZS, EY etc.) released white papers detailing adoption models and maturity frameworks (www.mckinsey.com) (www.ey.com). Thus emerged a consensus: to "scale AI effectively," a Pharma AI CoE with the right **design and governance** is indispensable (www.zs.com) (www.ey.com).

2. Organizational Design

Designing an effective Pharma AI CoE involves key choices about structure, governance, and resourcing. Here we examine common models, roles, and operating principles.

2.1 Organizational Models: Hub, Spoke, and Hybrid

Organizational archetypes. A seminal analysis by ZS Associates identifies three archetypal models for AI organization in pharma (www.zs.com) (table below). In practice, most companies start with either a centralized "hub" or a decentralized "spoke" model:

- **Hub (Centralized) Model:** All AI work is led by a central CoE team. Business units (R&D, manufacturing, commercial, etc.) submit requests to the CoE, which designs, builds, and governs solutions. Roughly *60%* of firms initially follow this approach (www.zs.com). Its strengths are in abundant funding, deep technical expertise, and unified leadership. The CoE can standardize processes rapidly and concentrate scarce AI talent. *However*, as ZS notes, such top-down models risk becoming disconnected from specific business needs (www.zs.com). The "vision at the top" may lose touch with local priorities, and business units may feel disempowered. Decision-making can bottleneck at the CoE.
- **Spoke (Decentralized) Model:** AI capability is embedded within each business unit. For instance, each drug development division, manufacturing site or marketing group may hire its own data science team. About *40%* of firms start here (www.zs.com). This ensures that AI use-cases are closely aligned with domain needs, and often fosters faster initial prototypes. *Yet*, it leads to duplication of effort and inconsistent tools: different teams may reinvent solutions that could be shared. It can be hard to enforce standards, and resourcing is uneven if one department has more budget or support.
- **Hub-and-Spoke (Hybrid) Model:** A compromise approach, the CoE (hub) provides strategy, governance and shared services, while smaller AI teams in business units (spokes) run use-case development (www.zs.com). The hub tackles enterprise-level tasks (data infrastructure, model inventories, compliance) while the spokes work on applications local to the unit. According to ZS, this hybrid model *"is more durable in the long run"*, leveraging the strengths of both centralized control and local agility (www.zs.com).

These models and their trade-offs can be summarized as follows:

Org Model	Structure	Advantages	Challenges
Centralized Hub	Single central AI team; business units submit requests.	Plentiful funding, deep expertise, unified vision (www.zs.com). Fast standardization of methods.	Can become disconnected from specific business needs (www.zs.com). Risk of bottleneck; slower responsiveness to local demands.
Decentralized Spoke	AI teams embedded in each business unit (R&D, supply chain, etc.).	Solutions closely aligned with local priorities; faster initial development (www.zs.com).	Duplication of work; inconsistent tools and data. Hard to share lessons across units (www.zs.com).
Hub-and-Spoke Hybrid	Central CoE sets strategy/governance (hub). Business-unit AI teams (spokes) build use-cases under that framework. (www.zs.com)	Balances oversight with agility. Allows reuse and scale of best-practices (www.zs.com).	Requires clear role definitions. Must manage tension between control and enablement (www.zs.com). Complex governance needed.

Table 1: Common AI organizational archetypes in pharma, with advantages and challenges. (Sources: ZS (www.zs.com)).

In practice, most mature programs evolve toward the hybrid hub-and-spoke model. Such a model retains centralized leadership and shared assets, yet empowers domain experts. *However*, firms must consciously address its pitfalls: distributing decision authority judiciously, clearly delineating responsibilities between the CoE and unit teams, and avoiding redundant governance layers (www.zs.com).

Example: Global Pharma CoE Initiative

As a real-world illustration, consider a large biopharmaceutical company that engaged consulting advisors to build a new AI CoE. The **challenge**, as documented in industry case studies, was fragmentation: siloed data science efforts across R&D, trials, and manufacturing, with no single strategy (peymancx.com). The approach was to create a *centralized AI Innovation Hub* that would drive collaboration across these functions. Key steps included:

- **Central Framework:** Designing a corporate AI roadmap and governance model. The CoE team defined standardized processes, from data preparation to model validation. (peymancx.com)
- **Use-Case Acceleration:** Deploying AI in discovery (predictive analytics for candidate molecules) and clinical trials (patient recruitment models) under CoE oversight (peymancx.com).
- **Manufacturing & Supply Chain:** Extending AI to optimize production and logistics (e.g. demand forecasting, predictive maintenance) through the hub's guidance (peymancx.com).
- **Governance and Compliance:** Aligning AI efforts with regulatory requirements – e.g. establishing validation protocols and security measures to build trust in AI decisions (peymancx.com).
- **Training and Culture:** Rolling out AI literacy programs across the organization to upskill employees and foster an innovation culture (peymancx.com).

Expected outcomes included faster drug discovery, streamlined trials, and improved operational efficiency (peymancx.com). This multifaceted initiative exemplifies how a hub structure can orchestrate company-wide AI adoption while still engaging each domain. (Further case studies are discussed in Section 5.)

2.2 Roles, Teams, and Talent

An effective AI CoE typically comprises a mix of technical and domain expertise. Key roles include:

- **CoE Director/Head:** Senior executive responsible for the AI strategy and alignment with business objectives. Often reports to CTO or CIO, and works closely with the C-suite to secure funding.
- **Data Scientists / ML Engineers:** Build and validate models. They require expertise in algorithms (classical ML and deep learning), as well as in working with biomedical data. According to McKinsey, teams must expand *"beyond traditional data science roles to include new skills – AI engineering, LLM fine-tuning, and business translation"* (www.mckinsey.com).
- **Data Engineers / Architects:** Develop the data infrastructure, pipelines and tools (e.g. ETL, data lakes, feature stores, MLOps) that make data accessible and models reproducible.
- **Domain Experts:** Scientists, clinicians, quality managers, regulatory affairs specialists, etc., who provide subject-matter knowledge. They help translate business problems into AI use-cases and validate the outputs of models. The hybrid CoE model relies on such experts at the spokes to guide AI development.
- **AI Ethics and Compliance Lead:** Ensures that AI projects adhere to ethical standards, privacy laws (HIPAA/GDPR), and industry regulations. This role might help formulate the company's AI ethics framework and liaise with legal/regulatory teams.
- **Change Manager/Program Manager:** Manages the AI transformation, prioritizes projects, and measures ROI. Also responsible for stakeholder communication and training programs.
- **IT/DevOps Specialists:** Maintain the computational infrastructure (cloud/on-prem), implement MLOps tools, and ensure security.

Smaller organizations might combine roles; large enterprises may form sub-teams (e.g. a Clinical AI team). A common pitfall is under-investing in *human* capabilities: surveys repeatedly highlight a shortage of AI-skilled professionals in pharma (www.ey.com). To address this, CoEs often run internal training and partner with universities to build the talent pipeline. Embedding data scientists within business units (spokes) also helps share domain knowledge – a practice encouraged by experts (www.mckinsey.com).

Illustrative Org Chart

A typical CoE organizational chart might show the CoE Director at the top, with sub-teams for *Data Platforms*, *Model Development*, *Governance/Quality*, and *Business Integration*. Each business unit in the company also fields its own AI/analytics lead, who coordinates with the CoE hub. While practices vary, one design principle is that the CoE should not simply be an "IT project team"; it needs authority and proximity to the business (often via dotted-line reporting or steering committees) to drive adoption (www.mckinsey.com) (www.deloitte.com).

2.3 Ways of Working

To spur innovation, many CoEs adopt agile and cross-functional ways of working. For example, forming small multidisciplinary squads for critical use cases (each with a data scientist, an engineer, and a domain expert) can accelerate pilots. However, each pilot should feed into the CoE's larger framework: models and code

should be documented in a central repository, and technology stacks should be standardized to avoid silos (www.mckinsey.com). Key practices include:

- **Use-Case Prioritization:** The CoE helps evaluate and prioritize AI use-cases based on business impact and feasibility. It ensures alignment with corporate strategy, avoiding the “AI for its own sake” trap (www.nature.com).
- **MLOps and Reuse:** Establishing an MLOps pipeline (model training, validation, deployment, monitoring) enables new projects to build on existing platforms. McKinsey notes a *platform-driven approach* provides reusable infrastructure for all use-cases (www.mckinsey.com).
- **Metrics and KPIs:** The CoE defines success metrics (e.g. reduction in time-to-market, accuracy of predictions, cost savings) and implements dashboards to track them. One industry mantra is *“measure the entire lifecycle”*, aligning AI performance with business outcomes (www.mckinsey.com).
- **Governance Procedures:** Formal processes (detailed in Section 3) are enforced through the CoE. This includes review checkpoints (e.g. design reviews, ethical clearance) for any AI project.

In summary, organizational design for a Pharma AI CoE involves balancing central strategy with distributed execution, hiring diverse talent, and instilling disciplined processes. The rest of this report will show that *governance* frameworks and *scaling playbooks* are equally critical to make these structures effective.

3. Governance: Managing Risk and Ethics

AI brings tremendous potential, but also unique risks. In healthcare and pharma, those risks translate into patient safety, data privacy, and overall trust. A Pharma AI CoE must embed robust governance to manage these issues across the **full AI lifecycle**. Effective governance ensures that AI tools do no harm and comply with laws and standards, while maintaining organizational accountability.

3.1 Core Principles of AI Governance

While many organizations have generic IT governance (for security, quality, financial oversight), AI governance requires additional layers. In essence, AI governance encompasses the policies, processes and oversight mechanisms specific to AI development and use. Key elements include:

- **Problem Framing:** Before building an AI system, governance should ensure there is a genuine, defined business problem to solve. Bodnari et al. emphasize that governance “must require a structured assessment of the underlying problem and a clear justification for why AI is the appropriate tool” (www.nature.com). This prevents “solutionitis” where AI is applied without need, which often leads to failure.
- **Data Governance:** Ensuring high-quality data is foundational. The CoE must work with data stewards to maintain data lineage, catalogues, and access controls. This includes patient and trial data. Assess obligations by jurisdiction, entity role, data status and intended use: the HIPAA Rules apply to covered entities and business associates (HHS); 21 CFR Part 11 concerns specified electronic records and signatures associated with FDA predicate rules and submissions (FDA). Review other privacy regimes, including GDPR, for applicability to the particular processing. Bodnari recommends *“integrating AI-*

specific data governance," such as standards for completeness, representativeness, and bias mitigation in training datasets (www.nature.com). For example, if an algorithm will predict disease risk, governance must check that training data appropriately reflects all patient demographics.

- **Model Review and Validation:** AI systems should be validated similarly to software, but with extra attention to their probabilistic nature. This often means performance testing, relevance of training data, and stress-testing for edge cases. The FDA's newest guidelines for AI/ML medical software (and concepts like *Good Machine Learning Practice*) highlight the need for continuous model monitoring and update protocols. Internally, governance often requires documenting model assumptions and conventions, and planning for periodic retraining.
- **Ethics and Fairness:** In pharma, biased algorithms could have grave consequences (e.g. unequal treatment recommendations). Governance must enforce fairness and transparency standards. Bodnari notes that governance mandates *"clear fairness standards, transparency, and accountability measures"* to reassure clinicians and patients (www.nature.com). Many companies form AI Ethics Committees or steering groups to evaluate new projects from a fairness and privacy perspective. For instance, an AI system used in patient triage might need to be reviewed as rigorously as a clinical trial protocol.
- **Accountability and Oversight:** The CoE should define roles responsible for AI outputs. For example, if an AI tool may influence a clinical decision, governance should specify who oversees that tool's use (a medical affairs exec, perhaps) and what recourse exists in case of error. This prevents the "black box" problem noted by Bodnari, where clinicians do not understand AI decisions (www.nature.com). Transparent audit trails (who approved what, when data/models were updated) are vital.

In summary, AI governance in a pharma context is about **embedding checks and balances** at every stage, to ensure systems are safe, ethical, and explainable. Trust cannot be assumed; it must be built by design. Bodnari et al. warn that without governance, *"AI systems can easily become sources of harm rather than benefit,"* and that *"governance is the bedrock of trust"* in healthcare AI (www.nature.com). Regulatory compliance (GxP, HIPAA, GDPR) thus becomes part of the CoE's remit.

3.2 Regulatory Considerations

Pharmaceutical companies operate under stringent regulations (FDA, EMA, PMDA, etc.) for drug development, manufacturing, and marketing. AI tools intersect with these regulations in various ways:

- **Clinical Decision Support and Diagnostics:** For AI tools regulated as medical devices, the CoE should determine the applicable route with regulatory affairs. In the United States, FDA pathways include 510(k) clearance, De Novo classification and premarket approval (PMA), depending on device classification and risk (FDA). In the EU, medical devices undergo conformity assessment before CE marking, with notified-body involvement where required; EMA provides specified scientific consultation and expert-panel support (EMA). Good Machine Learning Practice provides guiding principles for device development (FDA); FDA guidance recommendations do not themselves establish legally enforceable responsibilities (FDA). FDA's August 2025 guidance recommends that a predetermined change control plan (PCCP) describe planned modifications, the methods to develop, validate and implement them, and their impact assessment. Relief from additional marketing submissions applies only to modifications

specified in and implemented in accordance with the FDA-authorized PCCP, including its validation and implementation methods (FDA). The CoE should coordinate with regulatory affairs to align development and updates with applicable legal requirements.

- **Data Privacy:** Any AI using patient or citizen data must comply with privacy laws. This includes de-identification, consent management, and cross-border data rules. The CoE often sets policies to ensure datasets used for training (clinical trial data, EHRs) are handled correctly, and that AI outputs do not inadvertently reveal protected information.
- **Quality Systems:** Pharma is accustomed to Quality Management Systems (QMS) under GxP standards. A logical extension is to include AI development in the QMS framework. For example, code and models could be treated like controlled documents, requiring version control and change management. If an AI model is used to release batches, it may need formal validation akin to a laboratory method. McKinsey describes one organization specifying *observability and validation protocols in AI requirements from the start* (www.mckinsey.com).
- **Audit and Reporting:** Governance must ensure an audit trail for automated decisions. For instance, if an AI model triages patients, regulators may require manufacturers to report any adverse outcomes linked to the model. Early collaboration with legal and compliance teams (often part of the CoE's governance board) is crucial to meet both internal policies and external regulatory expectations.

In practice, many organizations form a cross-functional **AI governance committee** that includes Legal, Compliance, and Quality leads alongside data scientists. Its job is to align AI policies with external regulations and internal standards. These efforts address concerns that were highlighted by analysts: one EY report warns that *technical* challenges (privacy/security) and *ethical* concerns (bias/transparency) are key hindrances to AI adoption (www.ey.com). The CoE mitigates these by embedding data protection and algorithmic fairness into project lifecycles.

3.3 Frameworks and Lifecycles

Several published frameworks distill best practices for AI governance in healthcare. McKinsey recommends embedding risk management throughout the AI product lifecycle (McKinsey). Key governance checkpoints include:

1. **Project Initiation:** Determine whether AI is appropriate, define objectives and success metrics. Governance reviews project charters to ensure clear business justification (www.nature.com).
2. **Data Audit:** Assess and certify the data sources. Ensure representativeness and check for known biases according to domain requirements (www.nature.com).
3. **Model Development:** Follow secure coding and model training standards. Document design choices. Peer-review of model outputs may be required for sensitive use-cases.
4. **Validation:** Conduct testing on hold-out datasets or through retrospective trials. Regulatory-style "validation protocols" may be drawn up in advance (www.mckinsey.com).
5. **Deployment:** Before going live, have an approval meeting (like a Change Control board) that signs off on the model. Ensure audit logs will capture usage and outcomes.
6. **Monitoring:** Continuously track model performance, drift detection, and incidents. The CoE must have oversight of monitoring tools, triggering retraining or rollback if issues occur.

Throughout, emphasis is on *transparency and accountability*. For example, ethical guidelines might require that model decisions can be explained to clinicians and regulators. Bodnari specifically calls out the need for ongoing audits to catch unintended fairness problems (www.nature.com). Effective governance also means not treating risk as an obstacle, but as intrinsic to design: a leading practice is to build guardrails (technical and procedural) from day one (www.mckinsey.com).

In summary, governance is the **scaffolding** that holds the CoE's technical work together safely. It translates abstract principles (fairness, safety, compliance) into concrete rules and reviews. When done properly, it builds stakeholder trust: a physician or patient is far more likely to accept an AI recommendation if it comes buttressed by rigorous governance procedures (www.nature.com) (www.mckinsey.com). Conversely, the lack of clear governance is cited as a key reason many healthcare AI projects stall (mindsea.com) (www.ey.com).

4. Enterprise Scaling Playbook

A CoE's ultimate goal is not just to run pilots, but to **scale** AI so that it delivers measurable business value across the organization. This involves more than engineering work; it requires strategy, planning, and change management. Drawing on industry research, we outline key strategies and stages for moving AI from experiment to enterprise reality (a detailed checklist of actions is summarized in Table 2 below).

4.1 Key Strategies for Scaling AI

Recent thought leadership identifies several critical pillars for successful scaling:

- **Domain-Driven Focus:** Companies should align AI initiatives with core business domains (e.g. R&D, manufacturing, sales) rather than chasing isolated technical victories. McKinsey emphasizes a *"domain-driven approach"* where generative AI must reshape critical areas of the business (www.mckinsey.com). Survey data show that most life science firms prioritize R&D and commercial functions (28–38%) in their AI strategies (www.mckinsey.com). The point is clear: AI efforts should be embedded in value streams. There is "no such thing as a standalone gen AI strategy" – it must serve broader business objectives (www.mckinsey.com).
- **Organizational Transformation:** Scaling AI reshapes how the company operates. It requires changes in culture, processes, and talent. As noted earlier, McKinsey stresses that AI transformation is *"more than just tech"* (www.mckinsey.com). The operating model must adapt: decision rights may need redistributing (as in the hybrid hub-spoke model), and budgets may shift from traditional functions to AI initiatives. Talent strategies are paramount: besides hiring software engineers and data scientists, companies must develop new roles (AI engineers, MLops specialists) and "business translators" who can bridge technical and functional teams (www.mckinsey.com). One biotech example described launching an enterprise-wide upskilling program and appointing dedicated AI leadership in key functions, which significantly smoothed the scaling process (www.mckinsey.com).
- **Ecosystem and Partnerships:** No company can go it alone. The AI technology landscape evolves rapidly, so pharma must engage the broader ecosystem – including academic labs, tech partners, and startups (www.mckinsey.com). External collaborations (like NVIDIA co-innovation labs or partnerships

with AI/CDMO providers) allow access to cutting-edge tools and talent. McKinsey advocates cultivating a network of high-agility partnerships, and having “triggers” to scale up promising solutions (www.mckinsey.com). For example, once a pilot shows success, the CoE might leverage a cloud provider or contract research organization to industrialize the solution.

- **Platform and Infrastructure:** Reusable technical platforms are the foundation of scale. A *platform-driven approach* means building shared AI infrastructure early, so that new use cases plug into existing pipelines (www.mckinsey.com). The CoE should work on standardizing data ingestion, model training frameworks, and deployment pipelines. Standardization avoids repeated effort: it ensures that every new project benefits from the maturity of previous ones. McKinsey describes one unnamed life sciences company that spent three months defining a blueprint for insights and document platforms, enabling component reuse and rapid scaling across use cases (McKinsey).
- **Embedded Risk Management:** Scaling AI cannot come at the expense of safety. One McKinsey case study motto was *“slow down to speed up”*: companies formally defined validation, observability, and human-in-the-loop requirements before full rollout (www.mckinsey.com). In practice, this means at-scale models go through end-to-end risk assessment, often involving the CoE’s governance team. Risk management (data privacy, model biases, cybersecurity) is built in, not bolted on (www.mckinsey.com). A concrete best practice is to integrate AI governance checkpoints into the scaling lifecycle – for example, requiring a “fairness audit” sign-off before deployment.

These five strategies – domain focus, organizational transformation, ecosystem integration, platform scaling, and risk management – form the core of the enterprise playbook. Table 2 (below) summarizes these strategies and the associated actions.

Strategy	Key Actions
Domain-Driven Focus	Align AI projects with critical pharma domains (R&D, clinical, production, commercial) (www.mckinsey.com). Prioritize use cases that fundamentally reshape core processes. Ensure AI efforts directly support broader business objectives (e.g. higher success rates, faster trials). (www.mckinsey.com)
Org Transformation (Beyond Tech)	Rewire operating model: adapt roles, decision rights, and culture for AI. Launch company-wide upskilling (AI engineering, MLops, data literacy) (www.mckinsey.com). Hire or develop AI product owners/business translators. Set clear ROI/KPI metrics (e.g. time-to-market reduction, productivity gains) beforehand (www.mckinsey.com).
Ecosystem/Partnerships	Engage external partners (academia, tech vendors, AI-specialist CROs) to augment capabilities (www.mckinsey.com). Establish pilot-to-scale triggers and "sandbox" collaborations for agile innovation. Regularly update technology strategy via industry consortia and rapid prototyping with vendors.
Platform-Driven Approach	Build and expand common AI/ML platforms (data lakes, model repositories, MLOps pipelines) so new projects reuse infrastructure (www.mckinsey.com). Standardize data schema, APIs, and development tools across the enterprise. Implement “One Cloud/One Platform” strategy to consolidate AI assets.
Integrated Risk Management	Embed governance and compliance into the AI lifecycle. Involve risk & compliance teams early and continuously (www.mckinsey.com). Define clear ethical guidelines (for fairness, transparency). Implement monitoring dashboards for model drift/security. Prepare for evolving regulations (e.g., EU AI Act) by building required documentation and validation up front (www.mckinsey.com).

Table 2: Key strategies for enterprise-scale AI adoption in pharma (sources cited in text).

4.2 Phased Rollout and Pilots

Most organizations progress through stages of maturity (often called Foundational, Innovative, and Transformational (www.ey.com)). Early on (Foundational), the CoE may run pilots to demonstrate value and refine methodology. Each pilot should have a clear path for scaling:

1. **Inception and Incubation:** Identify a use case with executive sponsorship and measurable impact. Develop a proof-of-concept (PoC) in a limited scope, applying agile methods. For example, applying computer vision to automate QC in one production line.
2. **Validation and Build:** If the PoC shows results, expand to a Minimum Viable Product (MVP) for one site or one drug program. The CoE ensures the model is validated against real-world outcomes. Training on more diverse data may occur. Formal review by governance is conducted.
3. **Integration:** Deploy the model into routine operations at selected sites or business units. Provide training to users and integrate output into legacy systems (e.g. connecting the AI prediction to lab reporting tools). This often requires close collaboration between CoE tech teams and line managers.
4. **Scale-up:** Now Generalize the solution across geographies or additional products. Automate data pipelines for continuous input. The CoE measures business impact (e.g. X% acceleration in lab throughput) and monitors the system reliability. Knowledge from this scaling is fed back to improve the AI platform for the next use case (the “reuse cycle” mentioned by ZS (www.zs.com)).

A key insight is that *planning for scale begins at the pilot stage*. McKinsey recommends that even initial PoCs be “coded to scale” and use the eventual shared platform (www.mckinsey.com). Likewise, the CoE should define adoption criteria early (how many months to rollout, expected ROI) to prevent pilots floating aimlessly (www.mckinsey.com).

Finally, communication is crucial. The CoE should be transparent about both successes and failures. As the EY-Parthenon report notes, *ethical or technical failures* (e.g. data leaks or algorithmic errors) can undermine trust. Building stakeholder confidence often involves publishing “lessons learned” from early projects and updating governance policies accordingly.

4.3 Metrics and Value Tracking

Unlike conventional IT projects, AI initiatives can be unpredictable. Thus, the CoE should establish clear **metrics of value** from the outset. Examples in pharma include accelerated time-to-market for a candidate, improved success rate in clinical trials, reduced cost of goods in manufacturing, or higher patient engagement scores in a digital health program. McKinsey highlights the importance of defining these metrics *upfront*, e.g. time savings, productivity, or success probability improvements (www.mckinsey.com).

To track progress, organizations often create an AI dashboard overseen by the CoE. This may include technical metrics (model accuracy, data pipeline uptime) and business KPIs (cost savings, revenue impact). Regular reviews of these metrics (monthly/quarterly) ensure accountability. Over time, a consolidated “AI Scorecard” helps the C-suite gauge overall program health.

5. Case Studies and Examples

While comprehensive data on pharma AI CoEs is still emerging, several illustrative examples highlight different aspects of the playbook:

- **Global Pharma CoE (Peyman Case):** As described in Section 2, a leading pharma company built a central *AI Innovation Hub* with five focus areas (drug discovery, clinical trials, manufacturing, cross-functional standards, and governance) (peymancx.com). This approach embodies the hybrid model. They documented expected outcomes (speeding drug discovery, optimizing trials, reducing costs) (peymancx.com), and explicitly included “AI Governance & Compliance Strategy” as a pillar (peymancx.com). This underscores the integrated view: scaling AI required both technical deployment and structured oversight.
- **Startup Example (Formation Bio):** Formation Bio’s technical account describes developing its Model Context Protocol (MCP) gateway from March 2025 to make the MCP ecosystem more accessible to non-technical users (Formation Bio). ARK illustrates access and tooling infrastructure; this account should not be treated as evidence of model-training, validation or lifecycle-management capabilities.
- **PharmExec/ZS Interview (2026):** In a recent industry podcast, the CEO of ZS reported that for pharma to move from AI pilots to “scalable impact,” companies must shift from *informing* to *solving* patient needs, and adopt personalization at scale (www.pharmexec.com). He noted that regulatory fragmentation and a trust gap (among patients and providers) are the biggest barriers (www.pharmexec.com). This commentary reinforces the need for unified governance and patient-centric KPIs in any CoE strategy.
- **EY-Parthenon AI Maturity (2025):** EY and Microsoft’s report (launched at BioAsia 2025) introduced an AI maturity model with three stages (Foundational, Innovative, Transformational) (www.ey.com). While not a case study per se, the model provides a diagnostic tool for companies. For example, a firm might assess that it is *Foundational* in supply chain (just experimenting) but *Innovative* in drug safety. The report also quantified impact: *“75% of CXOs* surveyed in India’s life sciences industry said AI significantly cut costs and improved satisfaction” (www.ey.com), offering evidence that scaled AI can indeed deliver value. Such findings bolster the business case for establishing CoEs.
- **Academic/Industry Reviews:** Independent studies underscore the importance of governance and leadership. For instance, a scoping review on AI leadership in healthcare organizations (Sriharan et al, 2024) highlights that success depends on organizational readiness and Top-management support. Although not pharma-specific, it echoes the themes here: coordination, skills, and oversight are key. In a perspective article in *npj Digital Medicine*, Bodnari and Travis argue for governance frameworks to mitigate AI risks and build trust in healthcare organizations (Bodnari and Travis). This is a governance argument and set of recommendations, rather than a comparative study establishing better innovation outcomes for enterprises with robust governance.

Together, these cases and findings illustrate how CoEs act as the nexus for strategy, structure, and governance. Whether in a nascent biotech or a century-old pharma giant, the goal is the same: harness AI in a disciplined way to produce tangible benefits for patient health and business outcomes.

6. Implications and Future Directions

The journey of building and scaling a Pharma AI Center of Excellence has several broader implications and looks to the future:

- **Generative and Agentic AI:** The rise of large language models and multimodal AI opens new frontiers (e.g. rapid literature review, virtual assistants for clinicians). CoEs will increasingly focus on managing *generative AI* capabilities. This includes defining use-case boundaries (avoiding hallucination risk) and expanding infrastructure (GPU/TPU clusters). It also means updating governance: future models may require even more stringent validation and monitoring.
- **Regulatory Evolution:** Governments are actively shaping the AI landscape. The FDA's move to facilitate AI device updates (www.axios.com) and agencies adopting AI internally (FDA's own "Elsa" tool) signal that regulation is adapting, even as regulators themselves become savvy users. As of 26 October 2026, the European Commission's implementation FAQ sets out a phased AI Act schedule: prohibitions and AI-literacy provisions have applied since 2 February 2025, and governance and general-purpose AI obligations since 2 August 2025. High-risk AI rules start on 2 December 2027; rules for AI embedded in physical products, including medical devices, start on 2 August 2028. High-risk classification depends on intended purpose and the applicable legal categories; healthcare use alone does not automatically determine that status (European Commission). Pharma CoEs must stay ahead of these changes – for example, by preparing documentation and testing processes in anticipation (www.mckinsey.com).
- **Data Platforms and Privacy:** Data silos remain an obstacle, but progress in federated learning and secure multi-party computation could enable new use-cases (e.g. pharma-academic collaboration on patient data). CoEs will need to invest in data architecture that is both broad ("data lakes" for corporate data) and deep (interoperable with external databases). The emphasis on data lineage and governance will only intensify, as noted by governance experts (opengovasia.com).
- **Talent and Culture:** The talent gap in AI and data science is acute. CoEs will likely partner with universities, sponsor fellowships, and even acquire AI startups to secure skills. Internally, promoting an "AI-aware" culture is crucial: employees must trust and use AI tools rather than fear them. Communicating successes, offering training, and highlighting the augmentation (not replacement) of human roles will be ongoing tasks. A Deloitte forecast for 2030 underscores that bridging cross-functional gaps (tech vs life science) is key to overcoming talent constraints (www.deloitte.com).
- **Global and Ethical Considerations:** Pharma is global, so CoEs must handle multi-region issues. Data sovereignty (where data can be stored), regional regulations (e.g. China's AI policies), and linguistic differences (multilingual NLP models) all arise. Moreover, ethical considerations – such as equitable access to AI-derived treatments – may become larger societal questions. Centers of Excellence could expand their remit to include Responsible AI initiatives at scale, ensuring that innovations benefit all patient groups.
- **Monetization and ROI:** Ultimately, CoEs will be judged on outcomes. As the ZS CEO advised, pharma must move *beyond informing to solving* for patients (www.pharmexec.com). This means coalescing around patient-centric metrics (not just cost cutting). Successful CoEs will publish long-term impact –

e.g., demonstrating that AI-informed drug design led to faster approvals or that predictive maintenance saved millions in manufacturing downtime. Solid case studies and return-on-investment figures will reinforce the value of the CoE to skeptics (including boards and investors).

Conclusion

Building a Pharma AI Center of Excellence is a major strategic undertaking that touches every facet of the organization. The evidence from research and industry experience is clear: **structure and governance matter as much as technology**. A properly architected CoE – balancing central leadership with distributed innovation, and enforcing thorough governance – can turn isolated AI projects into a cohesive, scalable program.

Key takeaways include:

- **Organizational Design:** A hub-and-spoke model is generally optimal, with a central CoE aligning strategy and standards, and empowered business-unit teams executing use-cases (www.zs.com). Clear roles and funding mechanisms are essential to avoid duplication or bottlenecks.
- **Governance:** Trust is foundational. AI-specific governance frameworks (for fairness, data quality, validation, etc.) must be embedded throughout development (www.nature.com) (www.mckinsey.com). This includes complying with evolving regulations (e.g. FDA, EU AI Act) and documenting responsible practices from the start (www.mckinsey.com). Speaking broadly, those organizations that “crack the code” on AI governance will outpace those that don’t (mindsea.com).
- **Scaling Playbook:** Use-case selection should be business-driven (www.mckinsey.com). Scaling requires robust data and platform infrastructure (www.mckinsey.com), cross-functional skills (www.mckinsey.com), and partnerships (www.mckinsey.com). Metrics of success must be defined early (www.mckinsey.com), and progress must be measured systematically.
- **Culture and Change:** Beyond tech, success depends on leadership buy-in, clear communication, and talent development. Many companies underestimate the human factor. Upskilling existing staff and recruiting new skill sets (AI engineers, data stewards) is as important as installing servers.
- **Future Outlook:** The pace of AI innovation means the CoE must be forward-looking. It should experiment with new AI paradigms (e.g. agentic AI, digital twins), influence regulation through industry consortia, and continuously refresh its playbook.

Taken together, the organizational and governance elements described above constitute the “playbook” for scaling AI in pharma. Companies that undertake this enterprise transformation thoughtfully are likely to reap significant rewards: faster drug development, more efficient operations, and ultimately better patient outcomes. However, those that neglect governance or remain fragmented risk squandering AI’s promise. As pharma AI matures, the CoE will remain the heartbeat of this transformation – orchestrating innovation that is safe, compliant, and aligned with the mission of improving human health.

References: All claims and data above are substantiated by sources such as industry reports, academic articles, and case analyses (www.mckinsey.com) (www.zs.com) (www.nature.com) (www.axios.com) (www.ey.com), as cited in the text.

Source links

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

1. https://www.zs.com/insights/organizing-scaled-ai-pharma-hub-and-spoke-ai-model
2. https://www.mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry
3. https://www.nature.com/articles/s41746-025-01700-4
4. https://peymancx.com/ai-centre-excellence-pharma
5. https://www.ey.com/en_in/newsroom/2025/02/scaling-ai-in-pharma-ey-parthenon-microsoft-report-charts-path-to-enterprise-adoption-at-bio-asia-2025
6. https://www.pharmexec.com/view/scaling-ai-s-promise-in-healthcare-the-time-is-now
7. https://premainstream.com/2025/how-to-build-an-ai-center-of-excellence
8. https://hbr.org/2019/01/how-to-set-up-an-ai-center-of-excellence
9. https://www.novartis.com/stories/building-future-novartis-innovation-lab-ai-game-studio-xr-emerging-tech
10. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work
11. https://www.zartis.com/whitepaper/scaling-operations-in-ai-era
12. https://www.deloitte.com/middle-east/en/Industries/life-sciences-health-care/research/the-convergence-of-ai-technologies-and-human-expertise-in-pharma-r-and-d.html
13. https://www.hhs.gov/hipaa/for-professionals/covered-entities/index.html
14. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/part-11-electronic-records-electronic-signatures-scope-and-application
15. https://www.fda.gov/medical-devices/510k-clearances/medical-device-safety-and-510k-clearance-process
16. https://www.ema.europa.eu/en/human-regulatory-overview/medical-devices
17. https://www.fda.gov/medical-devices/artificial-intelligence-enabled-medical-devices/good-machine-learning-practice-medical-device-development-guiding-principles
18. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence
19. https://mindsea.com/ai-governance-in-healthcare-orgs
20. https://www.formation.bio/blog/ark-mcp-gateway
21. https://www.axios.com/2024/12/04/fda-ai-device-guidance
22. https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act
23. https://opengovasia.com/exclusive-the-end-to-end-ai-factory-building-trustworthy-and-scalable-ai-for-healthcare?c=vn

THE PUBLISHER

About IntuitionLabs

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

AI consulting and adoption

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

Software, data and life-science workflows

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

Enterprise platforms and regulated delivery

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

Work with IntuitionLabs

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

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

Website: https://intuitionlabs.ai

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