

AI Governance in Biotech: Council Charter and Decision Rights

6/30/2026 • 40 min read

ai governance biotech

biotech ai steering committee

ai council charter life sciences

ai decision rights framework

pharma ai governance

chief ai officer pharma

eu ai act pharma

fda ai guidance



Executive Summary

Artificial intelligence (AI) governance in biotech has moved from an ethics-committee side project to a board-level operating requirement. As of July 2026, **83%** of pharmaceutical executives report their organizations have formal governance committees for AI oversight, ahead of the **73%** reported by payer and provider peers, according to a Define Ventures survey of more than 40 executives across 15 of the top 20 global pharmaceutical companies (^[1] www.healthcarefinancenews.com). Yet the same period shows the boardroom lagging **operational adoption** industry-wide: only **27%** of corporate boards have formally added AI governance to committee charters even though **62%** hold regular AI discussions, per the National Association of Corporate Directors' 2025 survey (^[2] www.knostic.ai).

This report answers the central question facing biotech and pharmaceutical leadership: who owns AI, how should a steering committee or council be chartered, and where do decision rights sit when a model influences a **regulatory submission**, a manufacturing release, or a clinical trial design. The regulatory backdrop has hardened considerably. On **14 January 2026**, the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA) jointly published ten "Guiding Principles of Good AI Practice in Drug Development," the first transatlantic alignment of its kind, covering human-centric design, risk-based control, multidisciplinary expertise, data governance, and lifecycle management (^[3] www.fda.gov) (www.ema.europa.eu). Separately, the FDA's January 2025 draft guidance establishes a seven-step, risk-based credibility assessment framework for AI models used in regulatory submissions (^[4] www.fda.gov), and the European Union's AI Act reaches its principal compliance deadline of **2 August 2026** for high-risk AI systems under Annex III, with medical-device-embedded AI following in August 2027 (artificialintelligenceact.eu) (^[5] intuitionlabs.ai).

Organizationally, the market has settled on a recognizable pattern: a chartered AI governance council or steering committee, chaired by a senior executive (increasingly a **Chief AI Officer**), with defined decision rights that separate strategic sponsorship from day-to-day technical review. **Eli Lilly** appointed **Thomas Fuchs** as its first chief AI officer in October 2024 to set strategic direction across discovery, clinical trials, and manufacturing (^[6] www.biopharmadive.com); Pfizer added **Berta Rodriguez-Hervas** as chief AI and analytics officer in August 2024 (^[7] www.ciodive.com); and Lundbeck named **Markus Kede** as senior vice president and chief AI officer in March 2026, elevating him to the Executive Leadership Team (^[8] www.prnewswire.com). The commercial governance-tooling market reflects this urgency: Forrester forecasts spending on off-the-shelf AI governance software will more than quadruple to **\$15.8 billion** by 2030, while MarketsandMarkets projects the broader AI governance market growing from **\$0.89 billion** in 2024 to **\$5.78 billion** by 2029, a **45.3%** compound annual growth rate (^[9] www.forrester.com) (^[10] www.marketsandmarkets.com).

The report's central recommendation is that a defensible biotech AI governance council charter needs six recurring elements: a clearly bounded purpose and scope, membership with explicit quorum and voting rules, risk-tiered review triggers mapped to the FDA's context-of-use framework, a documented review and escalation process, formal interfaces with existing quality and data-governance bodies, and a scheduled revision mechanism. Companies that omit any of these six elements tend to produce governance structures that exist on paper but fail to resolve real disputes over model ownership, validation responsibility, or sign-off authority when an AI-influenced decision reaches a regulatory filing or a batch release. The remainder of this report details the regulatory drivers, the structural template, decision-rights allocation, quantitative benchmarks, and five real-world governance case studies from AstraZeneca, Eli Lilly, Sanofi, Novartis, and Lundbeck.

Introduction and Background

Biotech and pharmaceutical companies have adopted **artificial intelligence (AI)**, meaning computational systems that perform tasks associated with human intelligence such as pattern recognition, prediction, and generation, across nearly every stage of the value chain: target identification, [protein and molecule design](#), [clinical trial design](#) and recruitment, regulatory submission drafting, pharmacovigilance signal detection, and manufacturing quality control. **Generative AI (GenAI)**, the subset of AI that creates new content such as text, images, or molecular structures from learned patterns, has accelerated adoption further since 2023, with the FDA noting it has reviewed hundreds of drug and biologic submissions containing AI components since 2016 (^[11] [intuitionlabs.ai](#)).

That pace of adoption has outrun internal governance in most sectors, but biotech is a partial exception. A 2025 Define Ventures report based on interviews with more than 40 executives across 15 of the top 20 global pharmaceutical companies finds that **65%** of pharmaceutical executives identify AI as an immediate strategic priority, compared with **53%** of payer and provider executives (^[12] [www.healthcarefinancenews.com](#)). Governance structures have followed: pharma respondents report formal AI oversight committees at a higher rate (**83%**) than payers and providers (**73%**) (^[1] [www.healthcarefinancenews.com](#)). This is a sector-specific dynamic: unlike consumer software, an AI model that mis-selects a clinical trial endpoint, hallucinates a manufacturing deviation report, or biases a patient stratification algorithm carries direct patient-safety and regulatory-filing consequences, which pulls governance decisions into the same rooms as quality management, regulatory affairs, and pharmacovigilance.

Three converging forces make 2026 a decisive year for formalizing AI governance councils in biotech. First, regulatory clarity has arrived: the FDA-EMA joint principles of January 2026 establish a common transatlantic baseline, with the two agencies stating the principles “will underpin future AI guidance in the different jurisdictions and support enhanced international collaboration among regulators” ([www.ema.europa.eu](#)), while the EU AI Act’s high-risk obligations become enforceable on 2 August 2026 ([artificialintelligenceact.eu](#)). Second, board-level accountability structures remain immature relative to executive urgency: only **28%** of organizations report their CEO directly oversees AI governance, according to McKinsey’s 2025 “State of AI” survey, and just **17%** report board-level ownership (^[13] [www.knostic.ai](#)). Third, chief AI officer appointments at Eli Lilly, Pfizer, and Lundbeck signal that decision rights for AI are consolidating into named executive roles rather than diffuse working groups (^[14] [www.biopharmadive.com](#)) (^[8] [www.prnewswire.com](#)).

This report is written for biotech and pharmaceutical leaders, quality and regulatory affairs professionals, and technology executives who need a working answer to a deceptively simple question: who owns AI in our organization, and who decides what when a model touches a regulated process. It draws on regulatory guidance, peer-reviewed governance research, industry surveys, and named company case studies to build a practical council charter and decision-rights framework, anchored in conditions as of July 2026.

Who Owns AI in Biotech: Governance Structures and Reporting Lines

“Who owns AI” is not a single question in a biotech organization; it splits into ownership of strategy, ownership of technical validation, and ownership of risk acceptance, and conflating the three is the most common governance failure mode. A 2022 longitudinal case study of AstraZeneca’s ethics-based AI auditing process, conducted over 12 months by researchers from the Oxford Internet Institute working with AstraZeneca’s R&D Data Office, found that the central difficulty large multinational organizations face is “ensuring harmonised standards across decentralised organisations” and demarcating the scope of an audit across business units with different risk profiles (^[15] [www.arxiv.org](#)). AstraZeneca has published five ethical principles for data and AI since becoming, in its own account, the first global pharmaceutical business to publish such principles online: private

and secure, explainable and transparent, fair, accountable, and human-centric and socially beneficial (^[16] www.eversheds-sutherland.com).

Three ownership models recur across the industry:

- **Centralized Chief AI Officer (CAIO) model:** A single executive holds enterprise-wide accountability for AI strategy and, often, governance policy. Eli Lilly's Thomas Fuchs was tasked with setting "strategic direction" for AI initiatives spanning drug discovery, clinical trials, manufacturing, commercial activities, and internal functions (^[17] www.biopharmadive.com). Pfizer folded AI leadership into its digital organization under Berta Rodriguez-Hervas, reporting into the office of the chief digital and technology officer (^[7] www.ciodive.com). Lundbeck's Markus Kede joined the Executive Leadership Team reporting directly to the CEO, with an explicit mandate to "establish strong governance and responsible AI practices" (^[18] www.prnewswire.com).
- **Federated committee model:** Ownership is distributed across a cross-functional council with representation from R&D, quality, regulatory, legal, IT, and data science, and no single executive holds unilateral authority. A recommended approach for corporate AI policy formation calls for a "Steering Team" assembling legal and compliance, information security and IT, HR, business unit leaders, and a data scientist or AI ethics advisor before any policy is drafted (^[19] intuitionlabs.ai).
- **Embedded quality-system model:** AI governance is not a standalone function but an extension of the existing quality management system (QMS), following ISPE's GAMP 5 risk-based approach to computerized systems and ICH Q9 quality risk management (^[20] ispe.org).

In practice, most mid-size and large biotech organizations blend the second and third models, with a CAIO or equivalent chairing a council that reports jointly to the executive committee and interfaces formally with the QMS. The IAPP's 2025 AI Governance Profession Report, based on responses from more than 670 individuals across 45 countries, found that **50%** of AI governance professionals are typically assigned to ethics, compliance, privacy, or legal teams rather than to a dedicated AI function, and that organizations where privacy functions hold primary responsibility reported **67%** confidence in EU AI Act compliance, notably higher than other ownership models (^[21] iapp.org) (^[22] iapp.org).

Regulators themselves have adopted the committee model internally, which is instructive for sponsors deciding how much formality their own council needs: the FDA has coordinated its own internal AI governance through the establishment of the CDER AI Council within the Center for Drug Evaluation and Research, mirroring the cross-functional committee structure it now expects of sponsors (^[23] www.appliedclinicaltrials.com).

For adjacent advisors supporting this structural decision, the practical question is rarely "which org chart is correct" in the abstract, but which reporting line matches the company's existing quality culture and Veeva or clinical-systems architecture, since AI governance decisions increasingly need to be traceable through the same document-control and audit-trail infrastructure that already governs GxP records.

The Council Charter: Structure, Membership, and Purpose

A biotech AI governance council charter is the founding document that gives a governance body actual authority rather than advisory status. Drawing on recurring elements identified across multiple pharma AI governance charter templates and the ISPE/GAMP framework, a defensible charter contains six components.

Purpose and scope. The charter should state, in one paragraph, what decisions the council can make unilaterally, what it can only recommend, and what categories of AI system fall inside versus outside its authority (for example, explicitly excluding non-GxP productivity tools like generic writing assistants unless they touch

regulated data). Ambiguity here is the single most common charter defect: a charter that says the committee “oversees AI” without defining scope invites every department to either over-escalate trivial tools or under-escalate high-risk ones.

Membership and decision rules. Effective councils name specific seats, not just functions: a chair (often the CAIO or chief digital officer), and standing members from regulatory affairs, quality assurance, data science or R&D informatics, legal/compliance, information security, and a patient-safety or pharmacovigilance representative. The FDA-EMA joint principles explicitly require multidisciplinary expertise, stating that “multidisciplinary expertise covering both the AI technology and its context of use are integrated throughout the technology’s life cycle” ([24] www.fda.gov). Quorum rules and tie-breaking authority (usually resting with the chair) should be written explicitly rather than assumed.

Review triggers and risk tiering. The charter should map directly onto the FDA’s risk-based credibility framework, which assesses AI model risk based on **model influence** (how much the AI output drives the final decision without independent corroboration) and **decision consequence** (the severity of harm if the AI output is wrong), with higher-risk applications requiring more rigorous validation, extensive documentation, and lifecycle monitoring ([25] intuitionlabs.ai). The FDA-EMA principles state this directly: “the development and use of AI technologies follow a risk-based approach with proportionate validation, risk mitigation, and oversight based on the context of use and determined model risk” ([26] www.fda.gov). A tablet-coating-thickness prediction model and a clinical-trial-endpoint-selection model therefore warrant categorically different oversight intensity under this principle.

Review process. The charter should specify how a proposal reaches the council (intake form, sponsoring function, required documentation), what the council reviews (context of use, data lineage, validation evidence, bias assessment), and what outputs it produces (approval, conditional approval with monitoring requirements, or rejection with rationale).

Governance interfaces. The council cannot operate in isolation from the quality management system, data governance function, IT security, and, where applicable, the medical device or combination product regulatory pathway. ISPE guidance recommends integrating AI governance procedures with GAMP 5’s risk-based approach to compliant GxP computerized systems and ICH Q9 quality risk management processes for risk assessment, control, and review ([27] ispe.org). The International Council for Harmonisation has separately signaled that AI-based process models fall within its manufacturing-technology remit: an ICH reflection paper endorsed by the ICH Assembly on 8 October 2025 names “process modelling, including artificial intelligence (AI)-based models” as an example of an advanced manufacturing technology where “globally harmonised regulatory guidelines could facilitate the adoption and regulation” of the underlying technology ([28] admin.ich.org) ([29] admin.ich.org), giving governance councils a fourth standards body, alongside FDA, EMA, and ISO, to track for manufacturing-specific AI guidance.

Evolution mechanism. Given how fast both AI capability and regulatory expectation are moving, an annual or semi-annual charter review clause, with defined triggers for ad hoc revision (a new regulation, a governance incident, a major new AI capability), keeps the charter from becoming stale within eighteen months of adoption.

Table 1 below summarizes how the FDA-EMA joint principles map onto functional governance responsibilities inside a biotech organization, illustrating why a single-function committee (for example, an IT-only AI review board) cannot satisfy the full principle set.

Principle	Core Requirement (verbatim from the FDA-EMA document)	Primary Governance Owner
1. Human-Centric by Design	“The development and use of AI technologies align with ethical and human-centric values” ([30] www.fda.gov)	Business/clinical function owner
2. Risk-Based Approach	“Follow a risk-based approach with proportionate validation, risk mitigation, and oversight based on the context of use and	Quality/regulatory affairs

Principle	Core Requirement (verbatim from the FDA-EMA document)	Primary Governance Owner
	determined model risk" ([31] www.fda.gov)	
3. Adherence to Standards	"AI technologies adhere to relevant legal, ethical, technical, scientific, cybersecurity, and regulatory standards, including Good Practices (GxP)" ([32] www.fda.gov)	Quality assurance, IT security
4. Clear Context of Use	"AI technologies have a well-defined context of use (role and scope for why it is being used)" ([33] www.fda.gov)	Model owner/data science
5. Multidisciplinary Expertise	"Multidisciplinary expertise covering both the AI technology and its context of use are integrated throughout the technology's life cycle" ([24] www.fda.gov)	AI governance council (cross-functional)
6. Data Governance and Documentation	"Data source provenance, processing steps, and analytical decisions are documented in a detailed, traceable, and verifiable manner, in line with GxP requirements" ([34] www.fda.gov)	Data governance function
7. Model Design Practices	Development "leverages data that is fit-for-use, considering interpretability, explainability, and predictive performance" ([35] www.fda.gov)	Data science/model owner
8. Risk-Based Performance Assessment	Assessments "evaluate the complete system including human-AI interactions, using fit-for-use data and metrics appropriate for the intended context of use" ([36] www.fda.gov)	Regulatory affairs, clinical/quality
9. Lifecycle Management	"Risk-based quality management systems are implemented throughout the AI technologies' life cycles... undergo scheduled monitoring and periodic re-evaluation" ([37] www.fda.gov)	Quality management system owner
10. Clear, Essential Information	"Plain language is used to present clear, accessible, and contextually relevant information to the intended audience, including users and patients" ([38] www.fda.gov)	Pharmacovigilance/medical affairs

This mapping demonstrates that no single department can satisfy all ten principles alone; a defensible charter must formally assign, not merely imply, ownership of each principle to a named seat on the council, or the accountability gap becomes visible only after an incident or audit.

Decision Rights: Who Decides, Who Escalates, Who Delegates

A **RACI matrix** (Responsible, Accountable, Consulted, Informed), a tool that assigns each of these four roles to a stakeholder for a given decision or task, is the most common mechanism biotech companies use to formalize AI decision rights, though it has documented limitations. McKinsey's organizational research notes that RACI matrices "outline who owns which task so you know whom to reach out to should issues arise," but that many organizations find a strict RACI insufficient for genuinely cross-functional, ambiguous decisions and supplement it with explicit decision forums ([39] www.mckinsey.com).

In a biotech AI governance context, decision rights typically split into three tiers:

- **Decide (council or delegate authority):** Low-to-moderate risk AI applications, such as internal productivity tools or exploratory research models not feeding regulatory submissions, can be approved by a delegated sub-committee or even a single accountable owner (often the function head) without full council review, provided the intake criteria are met.

- **Escalate (full council review required):** Any AI system whose context of use touches a regulatory submission, a GxP-regulated process, patient safety, or personal health data crosses into mandatory full-council review under the FDA's risk-based framework, since these are precisely the applications where "model influence" and "decision consequence" are both high (^[40] intuitionlabs.ai).
- **Delegate (documented sub-authority):** Councils that try to review every AI use case personally become bottlenecks; mature governance structures delegate defined categories of low-risk decisions to business-unit AI leads or a "gatekeeper" role, while retaining audit rights and requiring quarterly reporting back to the full council.

A recurring organizing metaphor from Deloitte's public-sector AI governance research, cited in analysis of corporate AI policy design, describes three persona types useful for structuring decision rights: "Guides" who own policy and direction, "Guards" who enforce standardized quality-assurance checkpoints, and "Gadgeteers" who build and maintain the monitoring tools and feedback mechanisms that keep governance grounded in operational reality rather than paper compliance (^[41] intuitionlabs.ai). Applied to biotech, the "Guide" role typically sits with the CAIO or governance council chair, "Guard" functions map to quality assurance and regulatory affairs, and "Gadgeteer" responsibilities fall to data science and MLOps (machine learning operations, the practices that automate and simplify machine learning workflow deployment) teams responsible for drift monitoring and model registries (^[42] ispe.org).

Escalation pathways deserve particular attention because they are the most frequently under-specified element of biotech AI charters. A well-formed escalation clause names: (1) the trigger condition (a disagreement between the model owner and the reviewing quality function, a discovered bias or drift issue, or a near-miss safety event); (2) the escalation path (typically model owner to council chair to executive sponsor, with a defined maximum response time at each stage); and (3) the final decision authority, which in most charters rests with the council chair or, for the highest-risk disputes, the executive committee or chief medical/safety officer. Absent this specificity, disputes between, for example, a data science team confident in a model's performance and a regulatory affairs team concerned about explainability tend to stall indefinitely rather than resolve.

Regulatory and Standards Landscape Shaping Governance Design

Biotech AI governance charters do not exist in a vacuum; they are shaped directly by an increasingly dense and increasingly harmonized regulatory landscape spanning four major frameworks.

FDA guidance. The FDA's January 2025 draft guidance, "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products" (FDA-2024-D-4689), introduces a risk-based credibility assessment framework requiring sponsors to define the regulatory question and context of use (COU), assess model risk based on model influence and decision consequence, and execute a documented credibility assessment plan (^[43] intuitionlabs.ai). FDA Commissioner Robert Califf stated at the guidance's release that "with the appropriate safeguards in place, artificial intelligence has transformative potential to advance clinical research and accelerate medical product development to improve patient care," underscoring that the guidance is designed to enable rather than block adoption, provided credibility is established before submission (^[44] www.fda.gov).

FDA-EMA joint principles. The January 2026 joint publication "builds on collaborative work following the FDA-EU Bilateral meeting in April 2024," when both agencies began aligning on how AI was reshaping drug development (www.ema.europa.eu). Applied Clinical Trials Online, covering the announcement, quoted industry commentary that "regulators have been very clear in their positions and guidance on how to use AI in drug development more safely, effectively, and in line" with existing standards (^[45] www.appliedclinicaltrialsonline.com).

EMA reflection paper and Annex 22. Ahead of the joint principles, EMA's Committee for Medicinal Products for Human Use and Committee for Medicinal Products for Veterinary Use jointly adopted a "Reflection paper on the use of Artificial Intelligence in the medicinal product lifecycle" on 9 September 2024, following public consultation that closed in December 2023 (www.ema.europa.eu). EMA's Annex 22 is described as the first regulatory framework explicitly governing AI in drug production and, notably, prohibits generative AI for critical quality decisions in manufacturing even while other business functions are encouraged to adopt it (^[46] deepceutix.com). A governance council charter must therefore encode function-specific carve-outs, not a single blanket generative-AI policy.

EU AI Act. The Act, formally Regulation (EU) 2024/1689, entered into force on 1 August 2024 as the first comprehensive horizontal AI framework, applying a risk-based classification with prohibited, high-risk, limited-risk, and minimal-risk tiers (^[47] intuitionlabs.ai). AI systems embedded in devices regulated under the EU Medical Device Regulation (MDR) or In Vitro Diagnostic Regulation (IVDR) are automatically classified high-risk, described in the Commission's own summary as requiring compliance with "strict requirements, including risk-mitigation systems, high-quality of data sets, clear user information, human oversight" (commission.europa.eu), creating a dual compliance obligation where manufacturers must satisfy both AI Act and MDR/IVDR requirements simultaneously, though Article 8 allows integrating AI-specific testing and documentation into existing MDR conformity processes to avoid duplication (^[48] intuitionlabs.ai). High-risk AI systems under Annex III face a 24-month phase-in from entry into force, translating to full enforceability by 2 August 2026, though a November 2025 "Digital Omnibus" legislative proposal has signaled extensions for device-embedded AI to August 2027 (artificialintelligenceact.eu) (^[49] intuitionlabs.ai).

International ethics guidance. Two multilateral bodies supply the values-level language that many biotech responsible-AI principle documents draw on before a formal council charter translates them into operational controls. The World Health Organization's guidance on Ethics and Governance of Artificial Intelligence for Health, the product of eighteen months of deliberation among ethics, digital technology, law, and human rights experts, identifies six consensus principles intended to ensure AI "works to the public benefit of all countries" and holds both public and private stakeholders "accountable and responsive to the healthcare workers" and communities affected by its use (www.who.int). The Organisation for Economic Co-operation and Development's AI Principles, adopted in May 2019 and updated in May 2024, promote AI that is "innovative and trustworthy and that respects human rights and democratic values," organized around five values-based pillars: inclusive growth and well-being, human rights and democratic values, transparency and explainability, robustness and safety, and accountability, with 47 countries currently adhering to the recommendation (^[50] oecd.ai) (^[51] oecd.ai). Biotech governance councils operating across multiple jurisdictions frequently cite both frameworks in charter preambles as the shared ethical baseline underneath jurisdiction-specific binding rules such as the EU AI Act and FDA guidance.

Voluntary and standards frameworks. Outside binding regulation, two voluntary frameworks dominate governance design: the U.S. National Institute of Standards and Technology's AI Risk Management Framework (AI RMF), released 26 January 2023 and organized around four functions (Govern, Map, Measure, Manage), intended for voluntary use to incorporate trustworthiness into AI design, development, and evaluation (^[52] www.nist.gov), and ISO/IEC 42001, the international standard for an AI management system that ISPE recommends organizations reference when building formal AI governance policy documentation (^[53] ispe.org). A NIST generative AI profile followed on 26 July 2024, addressing risks unique to generative models (^[54] www.nist.gov). Governance councils increasingly use NIST AI RMF as the risk-taxonomy language and ISO 42001 as the auditable management-system backbone, layering GxP-specific requirements (ICH Q9 quality risk management, GAMP 5) on top (^[27] ispe.org).

For an adjacent advisory perspective, this layered regulatory picture is precisely why organizations building or extending Veeva Vault-based quality and regulatory systems increasingly need their AI governance charter to reference the same document-control and audit-trail architecture already in place for GxP records, rather than standing up a parallel, disconnected AI compliance system.

Company-level responsible AI principles as a governance input. Beyond binding regulation and voluntary standards, large biopharma companies have published their own responsible AI principles, and these documents function as de facto internal governance charters that a formal council later operationalizes. GSK's February 2024 public policy position on Responsible AI cites the World Economic Forum's prediction of "up to a 40 percent increase in labour productivity from AI in developed countries by 2035" as the strategic rationale for adoption, while explicitly committing to safeguard against "undermining human rights, data privacy and security," "errors that could cause patient harm," and "promoting discrimination and worsening inequalities" (^[55] www.gsk.com) (^[56] www.gsk.com). Novartis structures its ethical AI commitment around seven named themes, Empower Humanity, Accountability, Mitigate Bias, Respect Privacy, Transparent and Explainable, Safe and Secure, and Environmental Sustainability, applied across more than 100 documented AI use cases company-wide as part of a digital transformation effort underway since 2018 (^[57] www.novartis.com). For a governance council charter, the practical lesson from these company-level documents is that principle statements alone (fairness, transparency, accountability) are necessary but not sufficient; each of GSK's and Novartis's public commitments still requires the operational scaffolding, named seats, review triggers, escalation paths, described earlier in this report to become enforceable rather than aspirational.

Implementation Considerations and Process Changes

Standing up a functioning AI governance council is a multi-step organizational change, not a single charter-signing event, and most of the documented failure modes occur during implementation rather than design.

Step 1: Inventory before policy. A recommended sequence for corporate AI policy development begins with a comprehensive "AI asset register" audit of all AI and analytics tools already in use, since this step "often uncovers shadow AI, employees using external AI services without IT's knowledge" (^[58] intuitionlabs.ai). Shadow AI is a live and quantified problem: a 2025 industry benchmark found that **97%** of organizations that experienced AI-related security breaches lacked proper access controls, and **63%** of breached organizations lacked formal AI governance policies at all (^[59] www.knostic.ai) (^[60] www.knostic.ai).

Step 2: Regulatory mapping. Before drafting policy content, organizations must determine which laws and standards actually apply to their footprint, including data residency requirements, the EU AI Act if the organization markets any product in the EU (whose high-risk obligations bind "providers... regardless of whether they are based in the EU or a third country"), and FDA guidance for any AI touching regulatory submissions (artificialintelligenceact.eu).

Step 3: Draft with named roles, not templates. Copy-paste governance templates "often fail to fit a company's structure or culture" and must be tailored to the specific organization (^[61] intuitionlabs.ai). This is the point at which the six-element charter template described earlier should be adapted with actual named seats, actual escalation timelines, and actual quorum thresholds rather than generic placeholders.

Step 4: Build technical controls in parallel. Governance policy without technical enforcement is aspirational rather than operational. Recommended controls include an AI governance board or expanded risk committee, a mandatory risk-assessment workflow integrated into existing project management tooling, and technical controls such as secure development toolchains, access controls for AI cloud services, and model audit logging. This step is where the access-control gap documented above becomes most consequential, since it is precisely the absence of enforced technical controls, not the absence of written policy, that the breach data implicates most directly (^[62] www.knostic.ai).

Step 5: Schedule review, not just approval. A recommended practice is quarterly "policy health checks" by key stakeholders to address urgent regulatory or capability changes, in addition to a formal annual review cycle.

The gap between awareness and execution is well documented industry-wide, which suggests biotech organizations should treat council formation as a project with an executive sponsor and a defined timeline, rather than a standing agenda item that never quite gets prioritized. A 2025 AuditBoard research study found only **25%** of organizations across all sectors have a fully implemented AI governance program, attributing the gap to unclear ownership, limited expertise, and resource constraints rather than lack of awareness of the need ([63] www.knostic.ai). Gartner's 2025 poll of over 1,800 executive leaders found that **55%** of organizations have established a formal AI board or dedicated oversight committee, indicating the structural step is increasingly common even where full operationalization lags ([64] www.knostic.ai).

Data Analysis and Evidence

Quantitative signals from 2025 and 2026 surveys converge on a consistent picture: biotech and pharma governance adoption outpaces the broader enterprise average, but structural maturity within pharma itself remains uneven.

Table 2 below compares governance maturity indicators across the general enterprise population and pharmaceutical-specific survey data, drawn from independent research organizations.

Metric	General Enterprise (2025)	Pharmaceutical Sector (2025)
Formal AI governance committee in place	55% have an AI board or oversight committee ([64] www.knostic.ai)	83% of pharma executives report formal governance committees ([65] www.healthcarefinancenews.com)
AI as an immediate strategic priority	Not directly comparable	65% of pharma executives, vs. 53% of payer/provider peers ([12] www.healthcarefinancenews.com)
Fully implemented AI governance program	25% across all sectors ([63] www.knostic.ai)	Not separately reported; committee presence exceeds full implementation
Board has incorporated AI governance into committee charters	27% ([66] www.knostic.ai)	Not separately reported
CEO directly oversees AI governance	28% ([67] www.knostic.ai)	Not separately reported
Organizations actively working on AI governance	77% overall, rising to ~90% among AI users ([68] iapp.org)	Not separately reported

This comparison shows pharma outperforming general enterprise benchmarks on committee formation and strategic prioritization, but the underlying general-enterprise data on implementation gaps (governance committees existing without full operational maturity, and boards discussing AI without formal charter integration) should be read as a plausible cautionary signal for pharma as well, since sector-specific implementation-depth surveys were not identified in this research. Where the data diverges (a committee "in place" versus a governance program "fully implemented"), organizations should assume the more conservative, implementation-focused figure is the operationally meaningful one.

On the AI-in-drug-development side, the FDA-EMA joint principles' publication coincided with measurable clinical progress data: AI compresses early discovery timelines by an estimated **30-40%**, reducing preclinical candidate development from a traditional three-to-four years to **13-18 months**, and AI-discovered drugs are reportedly achieving **80-90%** Phase I success rates compared with a historical industry average of roughly **52%** ([69] deepceutix.com). No fully AI-developed drug had received FDA approval as of the January 2026 principles' publication, though industry commentary placed the probability of a first such approval in 2026-2027 at approximately 60%, an estimate that should be read as directional analyst opinion rather than an audited

regulatory figure. Given the pace at which those clinical claims are still being independently validated, governance councils should treat vendor- and analyst-sourced discovery-timeline figures as directional rather than as audited benchmarks.

On the governance-tooling investment side, two independent research firms provide converging but not identical forecasts. Forrester projects off-the-shelf AI governance software spending will more than quadruple by 2030, reaching **\$15.8 billion** and capturing **7%** of overall AI software spending (^[9] www.forrester.com). MarketsandMarkets projects the global AI governance market (encompassing platforms, data governance, MLOps tools, and consulting services) growing from **\$0.89 billion** in 2024 to **\$5.78 billion** by 2029, at a **45.3%** CAGR, with healthcare and life sciences named among the leading end-user segments alongside banking, financial services, and insurance (^[10] www.marketsandmarkets.com) (^[70] www.marketsandmarkets.com). These figures should not be read as interchangeable, since the two firms scope the market differently (Forrester scoping software spend specifically, MarketsandMarkets including services and consulting), but both point unambiguously toward a governance-tooling market growing far faster than overall AI software spending. A third estimate from Grand View Research, using a narrower market definition, values the global AI governance market at **\$308.3 million** in 2025, rising to an estimated **\$417.8 million** in 2026 and projected to reach **\$3,590.2 million** by 2033 at a **36.0%** CAGR, and separately forecasts the healthcare and life sciences vertical within that market will grow even faster, at a **39.9%** CAGR over the same period, reflecting the sector-specific urgency documented throughout this report (^[71] www.grandviewresearch.com) (^[72] www.grandviewresearch.com). That the three research firms disagree on absolute market size by more than an order of magnitude, while agreeing directionally on hyper-growth and on healthcare and life sciences as a leading vertical, is itself a useful governance signal: organizations budgeting for AI governance tooling should treat the CAGR trend as more reliable than any single absolute-dollar figure.

The NIST AI RMF's four core functions, Govern, Map, Measure, and Manage, provide the risk-taxonomy vocabulary increasingly used inside biotech governance councils even where GxP-specific frameworks like GAMP 5 and ICH Q9 supply the actual procedural detail (^[73] www.nist.gov). Under this shared vocabulary, "Govern" maps to the council charter and decision-rights structure itself, "Map" to the context-of-use documentation required by both FDA guidance and the FDA-EMA principles, "Measure" to the risk-based performance assessment and credibility evidence, and "Manage" to the lifecycle monitoring and drift detection obligations under Principle 9. Councils that adopt this four-function language when briefing boards and auditors report fewer translation gaps between technical AI teams and non-technical compliance reviewers, since NIST's taxonomy has become close to a lingua franca across regulated industries since its January 2023 release (^[74] www.nist.gov).

Case Studies and Real-World Examples

AstraZeneca: Ethics-Based Auditing as a Governance Operating Mechanism

AstraZeneca has arguably the most academically documented AI governance program in biopharma, the subject of a 12-month longitudinal case study by Oxford Internet Institute researchers working directly with the company's R&D Data Office (^[75] www.arxiv.org). The company describes itself as the first global pharmaceutical business to publish ethical AI principles online, organized around five pillars: private and secure, explainable and transparent, fair, accountable, and human-centric and socially beneficial (^[76] www.eversheds-sutherland.com). On its own sustainability disclosures, AstraZeneca states it applies "governance proportional to the impact and risk of our Data and AI Systems" and has implemented an AI Governance Framework, Risk Framework, and Playbook to operationalize these principles (^[77] www.astrazeneca.com) (^[78] www.astrazeneca.com). R&D Data Policy

Director Margi Sheth described the intent as “encouraging the right behaviors and culture” rather than a purely restrictive compliance exercise (^[79] www.eversheds-sutherland.com). The academic case study’s key finding for other biotech governance designers is that harmonizing standards across a decentralized multinational organization, and clearly demarcating audit scope, are the dominant operational challenges, not principle-writing itself (^[15] www.arxiv.org).

Eli Lilly: Consolidating Authority Under a Single Chief AI Officer

Eli Lilly announced on 9 October 2024 the appointment of Thomas Fuchs, formerly dean and department chair for AI and human health at Mount Sinai, as the company’s first chief AI officer, effective 21 October 2024 (^[6] www.biopharmadive.com). Lilly’s chief information and digital officer Diogo Rau framed the appointment around the scale of opportunity: “In this new era of technology, the potential for artificial intelligence and machine learning to revolutionize health care is immense” (^[80] www.biopharmadive.com). Fuchs’s mandate spans discovery, clinical trials, manufacturing, commercial activities, and internal functions, illustrating the enterprise-wide, cross-silo scope that a centralized CAIO model is designed to achieve (^[81] www.biopharmadive.com).

Sanofi: A Named Framework and Cross-Vendor AI Foundry

Sanofi has built what it calls the **RAISE** framework, a responsible AI framework emphasizing explainability, deployed alongside an internal “AI Foundry” partnership spanning three technology providers (AWS, Dataiku, and Snowflake) to ensure data lineage, bias monitoring, and traceability under what the company describes as a strict governance structure (zenn.dev). Sanofi CEO Paul Hudson has been an outspoken proponent of the strategy, telling McKinsey in a July 2025 interview that when the company recognized AI’s significance, “we were determined the company would not be a laggard. We would try a lot of things. We would make mistakes, and that was all acceptable as long as we were learning” (^[82] www.mckinsey.com). Separately, Sanofi Chief Digital Officer Emmanuel Frenehard described the company’s deliberate avoidance of open-ended pilots in favor of production-first deployment, stating the company approaches “AI developments as MVPs (minimum viable products), deploying them into production, learning from usage and user feedback, and then scaling across geographies” (^[83] enterpriseai.economictimes.indiatimes.com). This case illustrates a governance model where a named responsible-AI framework and a formal multi-vendor data-lineage architecture substitute for, or complement, a single centralized CAIO. Sanofi has extended this data-governance-first posture into 2026, selecting Snowflake as a core platform to run AI agents directly on governed enterprise data across research and development, procurement, IT, human resources, and field sales, rather than moving data between disconnected point applications (^[84] www.pharma-marketing.com).

GSK: Rebuilding the R&D Operating Model Around AI Governance

GSK has pursued a distinct governance path, framed less around a single named committee and more around rebuilding its research and development operating model so that governance is embedded in the platform itself. Trade press coverage of the company’s six-year transformation describes GSK moving from deploying isolated AI tools toward restructuring its “entire research and development operating model around AI, from data pipes in the basement to agents writing scientific hypotheses” (^[85] www.govinfosecurity.com). GSK’s public Responsible AI position, published in February 2024, frames this build-out explicitly around risk safeguards rather than capability alone, committing to guard against patient-harm errors and discriminatory outcomes even as the company pursues the productivity upside it associates with AI-driven R&D (^[86] www.gsk.com). GSK’s approach

illustrates that a governance council does not have to be the only locus of control; embedding governance requirements into the underlying data and platform architecture is a complementary strategy that reduces how much weight any single committee has to carry.

Moderna: Enterprise-Wide Deployment Under a Centralized Partnership

Moderna partnered with OpenAI beginning in early 2023 and has since deployed ChatGPT Enterprise to thousands of employees company-wide, extending the tool across legal, research, manufacturing, and commercial functions (^[87] [openai.com](#)) (^[88] [openai.com](#)). Moderna CEO Stéphane Bancel described the ambition in blunt terms: “We believe very profoundly at Moderna that ChatGPT and what OpenAI is doing is going to change the world” (^[89] [openai.com](#)). Within research data science specifically, Moderna’s approach is notably conservative on autonomy: the team, led by Eric Maher, deliberately avoids autonomous agents in favor of highly structured workflows, applies a “non-LLM filter” methodology that only reaches for generative AI after simpler Python or traditional machine-learning approaches have been exhausted, and matches evaluation rigor to consequence level for tasks like reformatting laboratory records into regulatory-compliant formats (^[90] [www.zenml.io](#)). This dual pattern, enterprise-wide access alongside function-specific risk-tiering, is a practical illustration of the risk-based governance principle in operation at working-team level.

Lundbeck: Governance Mandate Embedded in a New Executive Role

Lundbeck announced on 10 March 2026 the appointment of Markus Kede as senior vice president and chief AI officer, joining the Executive Leadership Team and reporting to President and CEO Charl van Zyl (^[91] [www.prnewswire.com](#)). Notably, Kede’s mandate is explicitly framed as a governance mandate rather than a purely technology mandate: he is tasked to “establish strong governance and responsible AI practices, build focused AI capabilities, and drive scalable deployment across functions” (^[92] [www.prnewswire.com](#)). Kede previously chaired Lundbeck’s US Digital & AI Strategy Committee, giving him direct prior experience with a federated committee model before taking on centralized authority (^[93] [www.prnewswire.com](#)). This progression, chairing a business-unit-level committee before assuming an enterprise CAIO role with an explicit governance mandate, is a pattern worth noting for mid-cap biotech organizations weighing how to sequence governance maturity as the organization scales.

Implications and Future Directions

Several structural shifts are likely to reshape biotech AI governance councils over the next 18 to 24 months. First, the FDA-EMA joint principles create pressure toward convergent, rather than duplicative, governance documentation: a company that builds separate US and EU AI governance processes will find increasing overlap wasteful once both agencies are working from a shared ten-principle baseline, and charter designers should build a single global framework with jurisdiction-specific annexes rather than parallel regional programs. As European Commissioner for Health and Animal Welfare Olivér Várhelyi put it, the guiding principles are “a first step of a renewed EU-US cooperation in the field of novel medical technologies,” suggesting joint agency documents will keep expanding rather than remain a one-time alignment exercise ([www.ema.europa.eu](#)).

Second, the role of Chief AI Officer appears to be consolidating from an emerging title into a standard C-suite fixture in large biopharma, following the pattern already visible in adjacent regulated industries; the trade press

noted the role “has proliferated across industries” with appointments at GE HealthCare, S&P Global, Boeing, and PwC preceding or coinciding with the pharma appointments profiled in this report (^[94] www.ciodive.com). Organizations without a named AI executive by 2027 risk being read by regulators, investors, and partners as governance laggards, independent of their actual technical maturity.

Third, international harmonization remains incomplete outside the FDA-EMA axis. The Japanese Pharmaceuticals and Medical Devices Agency (PMDA), China’s National Medical Products Administration (NMPA), and Australia’s Therapeutic Goods Administration (TGA) each maintain evolving, independent frameworks, meaning multinational biotech companies still face a jurisdictional patchwork the joint principles only partially resolve (^[95] deepceutix.com). Governance councils at globally operating biotechs should anticipate maintaining a jurisdiction-mapping annex to the core charter rather than assuming a single global standard will emerge in the near term.

Fourth, workforce constraints are likely to remain the binding limiter on governance execution speed, not technology or policy design. Principle 5 of the FDA-EMA framework mandates multidisciplinary teams, but “the talent pool is thin,” requiring professionals who understand both GMP deviation reports and gradient descent, a combination described as “scarce and expensive” (^[96] deepceutix.com). This scarcity reinforces the case for governance structures that route decisions efficiently to a small number of qualified reviewers rather than diffusing review responsibility broadly across under-trained staff, a pattern corroborated by the IAPP’s finding that **23.5%** of surveyed organizations cited finding qualified AI governance professionals as part of their delivery challenge (^[97] iapp.org).

Fifth, the persistent gap between board-level AI discussion (62%) and formal charter integration (27%) suggests governance formalization will continue for several more years before boards treat AI oversight as a standing, chartered responsibility comparable to audit or compensation committee oversight (^[2] www.knostic.ai). Biotech organizations that move early on formal charter language, rather than waiting for board-level pressure, are likely to face a smoother compliance path when EU AI Act enforcement intensifies after August 2026 and as FDA credibility-assessment expectations mature from draft to final guidance.

Advisory firms supporting life sciences organizations, including those specializing in Veeva-based quality and regulatory infrastructure, are increasingly asked to help translate these regulatory principles into system-level controls, such as ensuring an AI governance council’s decisions and model documentation are captured in the same audit-trail-compliant repositories already used for GxP change control, rather than existing as a separate, harder-to-audit governance silo.

Frequently Asked Questions (FAQs)

Who owns AI in a biotech organization?

Ownership typically splits across three dimensions: strategic ownership (often a Chief AI Officer, as at Eli Lilly, Pfizer, and Lundbeck), technical validation ownership (data science and R&D informatics, working under FDA’s risk-based credibility framework), and risk-acceptance ownership (the governance council chair or executive sponsor, for decisions crossing into GxP-regulated territory) (^[98] www.biopharmadive.com) (^[99] intuitionlabs.ai).

How do you set up an AI council in a biotech company?

Follow a five-step sequence: inventory existing AI use (including shadow AI), map applicable regulatory requirements (FDA, EMA, EU AI Act), draft a charter with named seats and explicit decision rules rather than a generic template (^[61] intuitionlabs.ai), build parallel technical controls (audit logging, access controls, model registries), and schedule recurring review cycles.

What should a biotech AI council charter include?

Six recurring elements appear across defensible charters: purpose and scope, membership and decision rules

(including quorum), risk-tiered review triggers mapped to context of use, a documented review process, formal interfaces with the quality management system, and a scheduled evolution mechanism.

What is a decision rights framework for AI, and how does RACI apply?

A decision rights framework specifies who can decide, who must be consulted, and who is merely informed for a given category of AI decision. RACI (Responsible, Accountable, Consulted, Informed) is the most common tool, though McKinsey research notes it is often insufficient alone for ambiguous cross-functional decisions and should be paired with explicit escalation pathways and a named final decision authority (^[100] www.mckinsey.com).

What does pharma AI governance require that general enterprise AI governance does not?

Pharma AI governance must map directly onto GxP quality systems (ICH Q9, GAMP 5), the FDA's context-of-use and credibility-assessment framework, and function-specific carve-outs such as EMA Annex 22's prohibition on generative AI for critical manufacturing quality decisions, none of which have general-enterprise equivalents (^[27] ispe.org).

When must biotech companies comply with EU AI Act high-risk requirements?

The principal deadline is 2 August 2026 for standalone high-risk AI systems under Annex III, with AI embedded in CE-marked medical devices following by August 2027 under the Digital Omnibus proposal's signaled extension (artificialintelligenceact.eu) (^[5] intuitionlabs.ai).

Is a Chief AI Officer necessary for biotech AI governance?

Not strictly, but it is increasingly common at large biopharma companies (Lilly, Pfizer, Lundbeck) as a way to consolidate accountability that would otherwise be diffused across a committee, and it can coexist with a federated council model, as at Lundbeck, where the CAIO chairs the council rather than replacing it (^[93] www.prnewswire.com).

Conclusion

AI governance in biotech has crossed a threshold in 2026: it is no longer a differentiator for a handful of leading companies but an operational expectation reinforced by joint FDA-EMA principles, an enforceable EU AI Act deadline, and a maturing chief AI officer role at major pharmaceutical companies. The organizations profiled in this report, AstraZeneca, Eli Lilly, Sanofi, Moderna, and Lundbeck, demonstrate that there is no single correct organizational design; centralized CAIO authority, federated cross-functional councils, and quality-system-embedded governance can all work, provided the charter is specific rather than aspirational.

The six-element charter template and the three-tier decision-rights structure (decide, escalate, delegate) presented in this report are not abstractions; they are drawn directly from patterns that recur across regulatory guidance, peer-reviewed governance research, and named company practice. The FDA's risk-based credibility framework, built on model influence and decision consequence, gives biotech organizations a ready-made calibration tool for review triggers, and the FDA-EMA joint principles give a common ten-point checklist against which any charter can be tested for completeness.

What remains unresolved, and what will likely define the next phase of biotech AI governance maturity through 2027 and beyond, is execution rather than design: closing the gap between the 83% of pharma executives who report a governance committee exists and the smaller, unmeasured fraction who can demonstrate that committee actually resolves disputes, tracks model lifecycle risk continuously, and satisfies a regulator's audit on demand. Biotech and pharmaceutical leaders who treat their AI governance charter as a living, tested operating document, reviewed at least annually and pressure-tested against real escalations, rather than a static compliance artifact, will be best positioned as FDA, EMA, and EU AI Act expectations continue to tighten through the remainder of this decade.

External Sources

- [1] <https://www.healthcarefinancenews.com/news/pharma-leading-payers-providers-ai-adoption#:~:While...>
- [2] <https://www.knostic.ai/blog/ai-governance-statistics#:~:While...>
- [3] <https://www.fda.gov/media/189581/download#:~:This%...>
- [4] <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological#:~:This%...>
- [5] <https://intuitionlabs.ai/articles/eu-ai-act-pharma-medical-device-compliance#:~:In%20...>
- [6] <https://www.biopharmadive.com/news/lilly-chief-ai-office-thomas-fuchs-mount-sinai/729326/#:~:Thoma...>
- [7] <https://www.ciodive.com/news/pfizer-chief-ai-officer-digital-leadership-team/724139/#:~:Pfize...>
- [8] <https://www.prnewswire.com/news-releases/lundbeck-appoints-chief-ai-officer-302709214.html#:~:H.%20...>
- [9] <https://www.forrester.com/blogs/ai-governance-software-spend-will-see-30-cagr-from-2024-to-2030/#:~:Forre...>
- [10] <https://www.marketsandmarkets.com/Market-Reports/ai-governance-market-176187291.html#:~:The%2...>
- [11] <https://intuitionlabs.ai/articles/fda-ai-drug-development-guidance#:~:The%2...>
- [12] <https://www.healthcarefinancenews.com/news/pharma-leading-payers-providers-ai-adoption#:~:only%...>
- [13] <https://www.knostic.ai/blog/ai-governance-statistics#:~:In%20...>
- [14] <https://www.biopharmadive.com/news/lilly-chief-ai-office-thomas-fuchs-mount-sinai/729326/#:~:Eli%2...>
- [15] <https://www.arxiv.org/pdf/2407.06232#:~:our%2...>
- [16] <https://www.eversheds-sutherland.com/en/luxembourg/insights/case-study-astrazeneca-featuring-margi-sheth-rd-dat-a-policy-director#:~:It%20...>
- [17] <https://www.biopharmadive.com/news/lilly-chief-ai-office-thomas-fuchs-mount-sinai/729326/#:~:He%20...>
- [18] <https://www.prnewswire.com/news-releases/lundbeck-appoints-chief-ai-officer-302709214.html#:~:In%20...>
- [19] <https://intuitionlabs.ai/articles/developing-corporate-ai-policy-governance-guide#:~:1.%20...>
- [20] <https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:The%2...>
- [21] <https://iapp.org/resources/article/ai-governance-profession-report#:~:the%2...>
- [22] <https://iapp.org/resources/article/ai-governance-profession-report#:~:When%...>
- [23] <https://www.appliedclinicaltrials.com/view/fda-ema-align-ten-principles-artificial-intelligence-use-drug-development#:~:the%2...>
- [24] <https://www.fda.gov/media/189581/download#:~:Multi...>
- [25] <https://intuitionlabs.ai/articles/fda-ai-drug-development-guidance#:~:asses...>
- [26] <https://www.fda.gov/media/189581/download#:~:The%2...>
- [27] <https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:Quali...>

- [28] https://admin.ich.org/sites/default/files/2026-03/ICH_ReflectionPaper_AdvancedManufacturing_2025_1008_0.pdf#:~:This%...
- [29] https://admin.ich.org/sites/default/files/2026-03/ICH_ReflectionPaper_AdvancedManufacturing_2025_1008_0.pdf#:~:Examp...
- [30] <https://www.fda.gov/media/189581/download#:~:Human...>
- [31] <https://www.fda.gov/media/189581/download#:~:follo...>
- [32] <https://www.fda.gov/media/189581/download#:~:AI%20...>
- [33] <https://www.fda.gov/media/189581/download#:~:AI%20...>
- [34] <https://www.fda.gov/media/189581/download#:~:Data%...>
- [35] <https://www.fda.gov/media/189581/download#:~:lever...>
- [36] <https://www.fda.gov/media/189581/download#:~:Risk%...>
- [37] <https://www.fda.gov/media/189581/download#:~:Risk%...>
- [38] <https://www.fda.gov/media/189581/download#:~:Plain...>
- [39] <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-organization-blog/the-limits-of-raci-and-a-better-way-to-make-decisions#:~:Many%...>
- [40] <https://intuitionlabs.ai/articles/fda-ai-drug-development-guidance#:~:high...>
- [41] <https://intuitionlabs.ai/articles/developing-corporate-ai-policy-governance-guide#:~:leadi...>
- [42] <https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:MLOps...>
- [43] <https://intuitionlabs.ai/articles/fda-ai-drug-development-guidance#:~:This%...>
- [44] <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions#:~:With%...>
- [45] <https://www.appliedclinicaltrials.com/view/fda-ema-align-ten-principles-artificial-intelligence-use-drug-development#:~:%E2%8...>
- [46] <https://deepceutix.com/insights/fda-ema-ai-principles#:~:EMA%2...>
- [47] <https://intuitionlabs.ai/articles/eu-ai-act-pharma-medical-device-compliance#:~:In%20...>
- [48] <https://intuitionlabs.ai/articles/eu-ai-act-pharma-medical-device-compliance#:~:Medic...>
- [49] <https://intuitionlabs.ai/articles/eu-ai-act-pharma-medical-device-compliance#:~:ai%20%...>
- [50] <https://oecd.ai/en/ai-principles#:~:The%2...>
- [51] <https://oecd.ai/en/ai-principles#:~:Today...>
- [52] <https://www.nist.gov/itl/ai-risk-management-framework#:~:The%2...>
- [53] <https://ispe.org/pharmaceutical-engineering/july-august-2024/artificial-intelligence-governance-gxp-environments#:~:stand...>
- [54] <https://www.nist.gov/itl/ai-risk-management-framework#:~:On%20...>
- [55] <https://www.gsk.com/media/m41d30ad/gsk-position-on-responsible-ai.pdf#:~:The%2...>
- [56] <https://www.gsk.com/media/m41d30ad/gsk-position-on-responsible-ai.pdf#:~:Under...>
- [57] <https://www.novartis.com/sites/novartiscom/files/novartis-responsible-use-of-ai-systems.pdf#:~:We%20...>

- [58] <https://intuitionlabs.ai/articles/developing-corporate-ai-policy-governance-guide#:~:This%...>
- [59] <https://www.knostic.ai/blog/ai-governance-statistics#:~:~:~:~:A%20s...>
- [60] <https://www.knostic.ai/blog/ai-governance-statistics#:~:~:~:~:AI%20...>
- [61] <https://intuitionlabs.ai/articles/developing-corporate-ai-policy-governance-guide#:~:~:~:~:each%...>
- [62] <https://www.knostic.ai/blog/ai-governance-statistics#:~:~:~:~:Only%...>
- [63] <https://www.knostic.ai/blog/ai-governance-statistics#:~:~:~:~:A%202...>
- [64] <https://www.knostic.ai/blog/ai-governance-statistics#:~:~:~:~:Accor...>
- [65] <https://www.healthcarefinancenews.com/news/pharma-leading-payers-providers-ai-adoption#:~:~:~:~:While...>
- [66] <https://www.knostic.ai/blog/ai-governance-statistics#:~:~:~:~:only%...>
- [67] <https://www.knostic.ai/blog/ai-governance-statistics#:~:~:~:~:only%...>
- [68] <https://iapp.org/resources/article/ai-governance-profession-report#:~:~:~:~:Of%20...>
- [69] <https://deepceutix.com/insights/fda-ema-ai-principles#:~:~:~:~:AI%20...>
- [70] <https://www.marketsandmarkets.com/Market-Reports/ai-governance-market-176187291.html#:~:~:~:~:The%2...>
- [71] <https://www.grandviewresearch.com/industry-analysis/ai-governance-market-report#:~:~:~:~:The%2...>
- [72] <https://www.grandviewresearch.com/industry-analysis/ai-governance-market-report#:~:~:~:~:The%2...>
- [73] <https://www.nist.gov/itl/ai-risk-management-framework#:~:~:~:~:Circu...>
- [74] <https://www.nist.gov/itl/ai-risk-management-framework#:~:~:~:~:Relea...>
- [75] <https://www.arxiv.org/pdf/2407.06232#:~:~:~:~:Over%...>
- [76] <https://www.eversheds-sutherland.com/en/luxembourg/insights/case-study-astrazeneca-featuring-margi-sheth-rd-dat-a-policy-director#:~:~:~:~:They%...>
- [77] <https://www.astrazeneca.com/sustainability/ethics-compliance/data-and-ai-ethics.html#:~:~:~:~:We%20...>
- [78] <https://www.astrazeneca.com/sustainability/ethics-compliance/data-and-ai-ethics.html#:~:~:~:~:We%20...>
- [79] <https://www.eversheds-sutherland.com/en/luxembourg/insights/case-study-astrazeneca-featuring-margi-sheth-rd-dat-a-policy-director#:~:~:~:~:We%20...>
- [80] <https://www.biopharmadive.com/news/lilly-chief-ai-office-thomas-fuchs-mount-sinai/729326/#:~:~:~:~:E2%8...>
- [81] <https://www.biopharmadive.com/news/lilly-chief-ai-office-thomas-fuchs-mount-sinai/729326/#:~:~:~:~:AI%20...>
- [82] <https://www.mckinsey.com/industries/life-sciences/our-insights/dont-delegate-the-ai-revolution-a-conversation-with-s-anofi-ceo-paul-hudson#:~:~:~:~:we%20...>
- [83] <https://enterpriseai.economictimes.indiatimes.com/news/case-studies/revolutionizing-drug-discovery-sanofis-ai-journey-from-snackable-tools-to-major-innovations/123778820#:~:~:~:~:We%20...>
- [84] <https://www.pharma-marketing.com/sanofi-picks-snowflake-to-scale-ai-across-drug-development/#:~:~:~:~:Sanof...>
- [85] <https://www.govinfosecurity.com/blogs/gsk-ai-driven-science-factory-p-4121#:~:~:~:~:The%2...>
- [86] <https://www.gsk.com/media/m41d30ad/gsk-position-on-responsible-ai.pdf#:~:~:~:~:This%...>
- [87] <https://openai.com/index/moderna/#:~:~:~:~:Moder...>
- [88] <https://openai.com/index/moderna/#:~:~:~:~:We%E2...>
- [89] <https://openai.com/index/moderna/#:~:~:~:~:We%20...>

- [90] <https://www.zenml.io/llmops-database/building-reliable-llm-workflows-in-biotech-research#:~:The%20...>
 - [91] <https://www.prnewswire.com/news-releases/lundbeck-appoints-chief-ai-officer-302709214.html#:~:H.%20...>
 - [92] <https://www.prnewswire.com/news-releases/lundbeck-appoints-chief-ai-officer-302709214.html#:~:He%20...>
 - [93] <https://www.prnewswire.com/news-releases/lundbeck-appoints-chief-ai-officer-302709214.html#:~:chair...>
 - [94] <https://www.ciodive.com/news/pfizer-chief-ai-officer-digital-leadership-team/724139/#:~:The%20...>
 - [95] <https://deepceutix.com/insights/fda-ema-ai-principles#:~:The%20...>
 - [96] <https://deepceutix.com/insights/fda-ema-ai-principles#:~:Princ...>
 - [97] <https://iapp.org/resources/article/ai-governance-profession-report#:~:Of%20...>
 - [98] <https://www.biopharmadive.com/news/lilly-chief-ai-office-thomas-fuchs-mount-sinai/729326/#:~:He%20...>
 - [99] <https://intuitionlabs.ai/articles/fda-ai-drug-development-guidance#:~:asses...>
 - [100] <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-organization-blog/the-limits-of-raci-and-a-better-way-to-make-decisions#:~:Many%...>
-

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.

© 2025 IntuitionLabs.ai. All rights reserved.