



INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Enterprise AI Governance in Pharma: GxP & Compliance

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 · 2026-02-08 · ai governance · pharmaceutical industry · gxp compliance · fda regulations · data integrity · machine learning validation · risk management · gamp 5

Original article: <https://intuitionlabs.ai/articles/pharma-ai-governance-gxp-compliance>



In this article

1. Executive Summary
2. Introduction and Background
3. Regulatory Environment and Guidelines
4. Core Governance Domains and Practices
5. Comparative Table of Governance Domains
6. Case Studies and Real-World Examples
7. Implications and Future Directions
8. Conclusion

Executive Summary

Artificial intelligence (AI) promises transformative benefits across the pharmaceutical industry—from accelerating drug discovery and development to optimizing manufacturing and patient safety—while also presenting unique challenges in highly regulated environments. Effective **enterprise AI governance** in regulated pharmaceutical companies must balance innovation with robust oversight. This involves integrating AI into existing quality management (GxP) frameworks, embedding risk-based controls, ensuring data integrity and patient privacy, and aligning with evolving regulations and ethical standards (www.frontiersin.org) (www.auriacompliance.com). The corporate governance perspective stresses the need for clear accountability, cross-functional collaboration, and ongoing education to “bridge the gap between principles and practice” in AI deployment (pmc.ncbi.nlm.nih.gov) (www.frontiersin.org).

Recent developments highlight this urgency. The U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA) have issued new guidance and principles to govern AI use in drug and device development (www.fda.gov) (www.ema.europa.eu). For example, FDA’s 2025 draft guidance introduces a risk-based framework for assessing AI model credibility in drug submissions (www.fda.gov), while the FDA/Health Canada/MHRA consortium published “Transparency for ML-Enabled Medical Devices” (2024), emphasizing human-centered, lifecycle transparency in AI medical devices (www.fda.gov) (www.fda.gov). EMA adopted its reflection paper in September 2024 after the 2023 consultation and, jointly with FDA, published ten good AI practice principles in January 2026 covering the medicines lifecycle (www.ema.europa.eu) (www.ema.europa.eu). Globally, regulators (including MHRA, NMPA) and standards bodies (NIST, ISO, OECD, WHO) are converging on frameworks for ethical, accountable AI use in healthcare. Key principles emerging from these efforts include rigorous risk assessment, model validation, data governance (privacy and integrity), transparency, and human oversight (www.frontiersin.org) (www.gov.uk).

In practice, pharmaceutical AI governance must be **pragmatic and integrated**. Companies should build on existing quality and compliance systems (e.g., GAMP 5 computerized system validation, 21 CFR Part 11) rather than creating isolated processes (www.frontiersin.org) (pmc.ncbi.nlm.nih.gov). For example, AstraZeneca’s AI governance experience emphasizes harmonizing AI risk assessments with traditional quality processes, using risk tiers (low/medium/high) to scale controls, and leveraging familiar structures for ethics and change management (www.frontiersin.org) (www.frontiersin.org). Adoption of AI requires cross-functional oversight—combining R&D, quality, IT, regulatory affairs, and business stakeholders—with support from senior leadership and possibly dedicated roles (e.g. an AI Officer or Ethics Board) (www.frontiersin.org). Training and documentation (e.g. model cards, audit trails) support traceable governance (www.frontiersin.org) (www.fda.gov). The [adopted EMA paper](#) and [January 2026 EMA–FDA principles](#) distinguish the medicines context; [NIST AI RMF](#) supports multidisciplinary risk management.

This report provides a comprehensive analysis of AI governance in regulated pharma, drawing on regulatory documents, industry frameworks, and case studies. It covers the historical and current landscape of AI in pharma, detailed governance domains (data, validation, ethics, etc.), implementation best practices, and future outlook. Throughout, we highlight specific data and guidelines, such as the fact that by 2020 *90% of large pharma have initiated AI/ML projects* (ispe.org), and that FDA has reviewed over *500 submissions with AI components since 2016* (www.fda.gov). Finally, we look ahead to emerging trends (AI Act, generative AI, global convergence) and their implications for regulated companies.

By synthesizing expert opinions and evidence, this report concludes that rigorous, **enterprise-wide AI governance** is not merely a compliance exercise but a strategic imperative for pharmaceutical companies to safely harness AI's potential while protecting patients, data, and company integrity. We present detailed recommendations—rooted in current guidance and case experience—for building such governance frameworks, with governance “guardrails” that turn “Can we use this?” into “Yes, safely” (pharmaconsulting.ai). Regulatory statements are accompanied by source links; recommendations and illustrative scenarios should be read as proposed governance practices.

Introduction and Background

The pharmaceutical industry stands at the cusp of an AI-driven transformation. Advances in machine learning (ML) and generative AI have opened new possibilities in drug discovery, clinical trials, manufacturing, and more. For example, deep learning models like AlphaFold2 and ESMFold can predict protein structures, radically accelerating early-stage research (www.mckinsey.com). Pharmaceutical companies report rapidly increasing AI adoption; by 2020, *90% of large pharma* had begun AI/ML projects (ispe.org). AI is being applied across the product lifecycle: to analyze real-world data, optimize trial design, automate documentation, detect safety signals, and personalize treatments. McKinsey estimates that generative AI could generate \$60–110 billion in annual potential economic value for the pharma and medical-product industries (www.mckinsey.com). Major firms (e.g., Merck, Novartis, Roche, AstraZeneca, Pfizer) are investing heavily in AI R&D, including partnering with AI startups and academic labs.

However, AI in pharma brings unique risks and regulatory complexities. Pharmaceuticals are among the most heavily regulated industries, with strict controls for quality, safety, and data integrity (GxP). Assess each AI system by entity, jurisdiction, intended use, and record type to identify applicable legal requirements and relevant guidance, including FDA's **21 CFR Parts 210/211 (GMP)**, **Part 11 (electronic records/signatures)**, ICH Q7/Q8, and Good Clinical/Pharmacovigilance Practices (GCP/GVP). Binding regulations, nonbinding guidance, and voluntary industry practices have different scopes and should be distinguished in the applicability assessment. AI's dynamic nature (e.g. continuous learning, opaque “black boxes”) challenges these static compliance models. Without clear governance, AI-driven processes risk introducing undocumented decisions, data breaches, or biased outputs that could harm patients or trigger non-compliance (pmc.ncbi.nlm.nih.gov).

Data integrity is foundational in pharma; see [MHRA's GxP data integrity guidance](#). Regulator guidance emphasizes that data must be **Attributable, Legible, Contemporaneous, Original, and Accurate (ALCOA)**, plus Complete, Consistent, Enduring, and Available (ALCOA+) (www.auriacompliance.com). AI systems must generate data and decisions that meet these criteria (e.g. logs, audit trails). For instance, using patient data to train models calls for an assessment of applicable privacy law and protections against unauthorized disclosure; HIPAA applies to [covered entities and business associates](#) (talkingmedicines.com) (www.auriacompliance.com). Indeed, AI can improve data integrity by detecting anomalies, but it also raises cybersecurity and validation challenges (www.auriacompliance.com) (www.auriacompliance.com). The integration of AI into GxP domains thus demands a robust governance layer that ensures AI outputs are traceable, validated, and respect regulatory and ethical boundaries.

Historically, pharma has navigated similar transitions whenever new technologies emerged (computerized systems, bioinformatics, e-records). The mantra “AI is different” is often used, but in many ways existing GxP principles still apply. Industry experts argue that **operationalizing AI governance** means building on current governance frameworks, not inventing them anew (www.frontiersin.org) (ispe.org). For example, the ISPE’s GAMP™5 guidelines, while written before the AI boom, provide a risk-based approach to computerized system validation that can extend to AI/ML systems (pmc.ncbi.nlm.nih.gov) (pmc.ncbi.nlm.nih.gov). For electronic records, assess **21 CFR Part 11 applicability** record by record against the underlying predicate rules. FDA recommends documenting that determination; a record used to generate a submission is not itself a Part 11 record unless it is otherwise required under a predicate rule and maintained electronically. In practice, pharmaceutical companies must integrate specific AI controls (e.g. bias testing, model monitoring) into these existing structures.

Recent global events underscore the stakes. In 2023–2025, regulators have accelerated AI policy. The 2023 Executive Order 14110 is historical: **NIST records its rescission on 20 January 2025**. The FDA released AI/ML guidance for medical devices and (Jan 2025) a draft guidance on AI in drug development (www.fda.gov) (www.fda.gov). EMA published a draft reflection paper in July 2023, closed consultation on 31 December 2023, and published the adopted paper on 30 September 2024 (www.ema.europa.eu) and, with FDA, published ten AI principles (January 2026) for medicine development (www.ema.europa.eu). The UK’s MHRA unveiled a five-principle AI strategy (2024) emphasizing safety, transparency, fairness, accountability, and redress (www.gov.uk). The **EU AI Act (Regulation (EU) 2024/1689)** entered into force on 1 August 2024 and has phased application dates. In sum, regulatory and societal expectations for AI governance have dramatically intensified.

Despite this, many companies still lack mature AI governance. Common barriers include unclear ownership of AI risks, generic policies (not tailored to specific departments), confusion over when validation is needed, and skills gaps among staff (pharmaconsulting.ai). For example, a pharma team might not know if clinical trial recruitment AI falls under IT or QA oversight, or how to document model decisions in a way acceptable to auditors. This report will chart a path through these complexities. We first survey the current landscape of guidelines and frameworks (Table 1 below summarizes key initiatives by region). We then examine core governance domains (data/privacy; model development and validation; risk management; ethical oversight; organizational roles; generative vs traditional AI, etc.), illustrating each with data and case examples. Throughout, we integrate expert recommendations and argue for embedding AI governance into the enterprise risk and quality systems of regulated pharma companies. Tables 1 and 2 encapsulate some of these comparisons and considerations.

Table 1. Key regulatory and industry frameworks for AI governance in pharmaceutical and healthcare sectors.

Region/Body	Framework/Guidance	Scope/Highlights
USA (FDA)	<i>FDA AI/ML Drug/Biologics Guidance (2025, draft)</i> – Risk-based framework for AI model credibility in drug submissions; emphasizes “context of use” for each model (www.fda.gov).	AI in drug dev & submissions; model validation and trust.

Region/Body	Framework/Guidance	Scope/Highlights
USA (FDA)	<i>Good Machine Learning Practice (GMLP) Guiding Principles (FDA/Health Canada/MHRA, 2021)</i> – Baseline principles (e.g. quality, transparency) for AI medical devices.	Foundational best practices for AI/ML devices.
USA (FDA)	<i>ML Device Transparency Guiding Principles (2024)</i> – Joint FDA/Health Canada/MHRA document; outlines “who, what, where, when, why, how” of ML system information; stresses human-centered design and comprehensive communication to users (www.fda.gov) (www.fda.gov).	Communication requirements for ML-based devices.
USA (NIST)	<i>AI Risk Management Framework (RMF v1.0, 2023) & Generative AI Profile (2024)</i> – Voluntary cross-sector guidelines to incorporate trustworthiness (fairness, explainability, security) into AI life cycle (www.nist.gov).	Enterprise risk management model for AI; includes GenAI-specific profile.
EU (EMA/FDA)	<i>Good AI Practice Principles (EMA–FDA, January 2026)</i> – Ten joint principles for AI in medicine R&D and lifecycle, from discovery through manufacturing to post-market safety (www.ema.europa.eu).	High-level principles; assess applicable law separately.
EU (EMA)	<i>Reflection paper on AI in the medicinal product lifecycle (adopted 2024)</i> – Adopted EMA reflection paper ; emphasizes human-centric AI, legal/ethical compliance, early regulator interaction if AI affects benefit-risk (www.ema.europa.eu) (www.ema.europa.eu).	Outlines thinking on AI use cases (discovery, trials, PV) and regulatory challenges.
EU (EMA)	<i>Network Data Steering Group (NDSG) workplan 2026–2028</i> – Action plan for guidance on AI tools, frameworks, and capacity-building in EU medicines regulation (www.ema.europa.eu).	Future EMA guidance development on AI in pharma.
EU (EC)	<i>EU AI Act — Regulation (EU) 2024/1689</i> — Enacted; entered into force on 1 August 2024.	Phased application; classify each use by legal criteria. See Commission timeline and classification guidance .
UK (MHRA)	<i>AI Regulatory Strategy (2024)</i> – MHRA sets five strategic principles: Safety/Security, Transparency/Explainability, Fairness/Accountability, Contestability/Redress; positions AI as critical in future healthcare regulation (www.gov.uk).	Covers both AI as medical products (AIaMD) and AI as regulator of products; risk-proportionate controls and innovation.
International (WHO/Others)	<i>WHO Guidance on Ethics and Governance of AI for Health</i> – (2021) Global guidance emphasizing equity, privacy, and human rights in health AI.	High-level ethical framework; not industry-specific.
Industry Standards	<i>ISPE GAMP 5 (Second Edition), Appendix D11</i> — ML subsystem lifecycle guidance; ISPE explanation . The GAMP Artificial Intelligence Guide was published in July 2025.	Industry good-practice guidance used alongside applicable law; the AI Guide complements GAMP 5 Second Edition.

Table 1: A snapshot of regulatory and framework guidance influencing AI governance in pharmaceuticals. Details are drawn from official regulatory releases and industry publications cited above.

Regulatory Environment and Guidelines

Pharmaceutical companies must navigate a complex regulatory landscape when deploying AI. Unlike non-regulated sectors, pharma AI initiatives must integrate with laws governing drugs, devices, patient data, and clinical trials. Key jurisdictions (US, EU, UK, others) are actively publishing guidelines.

United States (FDA, ONC, etc.) – The FDA has long regulated AI as part of “Software as a Medical Device” (SaMD). In January 2025, FDA issued **Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations**, a draft covering AI-enabled device software functions and marketing-submission documentation across the device lifecycle. It contains nonbinding recommendations and is not for implementation. Separately, in 2024, FDA (with international partners) published the *Transparency for ML-Enabled Medical Devices: Guiding Principles* (www.fda.gov) (www.fda.gov). This document instructs manufacturers to communicate AI/ML device information (purpose, performance, limitations) clearly to providers/patients, using user-centered design. Separately, HHS/ONC’s **HTI-1 final rule** updates the ONC Health IT Certification Program and establishes transparency requirements for AI and other predictive algorithms in certified health IT. Assess that scope separately from FDA device regulation and drug submissions.

For drug development, the landmark event was FDA’s January 6, 2025 announcement of draft guidance **Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products** (www.fda.gov). This guidance proposes a **risk-based framework** to assess the *credibility* of AI models in regulatory submissions. As a draft with nonbinding recommendations, it is not for implementation. It recommends that sponsors clearly define the “context of use” for each AI model (e.g. predicting trial outcomes) and provide evidence (validations, testing plans) commensurate with the risk. Importantly, FDA encourages *early engagement*: companies are urged to discuss AI plans with FDA review teams well before submission.

FDA’s Office of the Chief Scientist and CDER have coordinated these efforts, establishing a CDER AI Council in 2024 to oversee policy consistent with the agency’s strategic priorities (www.fda.gov). FDA documents note that since 2016, over *500 drug and biologic submissions contained AI components*, and the rate is accelerating (particularly in oncology and neurology) (www.statnews.com) (www.fda.gov). These statistics underscore that AI is no longer hypothetical but a routine element of modern submissions, requiring clear rules.

Europe (EMA and National Regulators) – EMA issued a draft reflection paper in July 2023; its consultation closed on 31 December 2023. The **adopted reflection paper** was published on 30 September 2024 and addresses AI across the medicinal product lifecycle. It should be read in the context of applicable EU law.

In January 2026, EMA and FDA published **ten guiding principles of good AI practice in drug development**. They call for proportionate validation, risk mitigation and oversight based on context of use and model risk. EMA’s **AI overview** now identifies the NDSG workplan for 2026–2028 in its AI observatory section.

The **EU AI Act (Regulation (EU) 2024/1689)** entered into force on 1 August 2024. Prohibited practices and AI-literacy obligations applied from 2 February 2025; governance and general-purpose AI obligations applied from 2 August 2025. The Commission identifies general application from 2 August 2026, Annex III high-risk rules from 2 December 2027, and Annex I regulated-product high-risk rules from 2 August 2028. These later deadlines do not defer obligations already applicable.

Determine high-risk status by intended purpose and the **relevant legal classification criteria**, rather than assuming every clinical or quality-control application qualifies. The product route covers AI products or safety components within the relevant sectoral legislation when subject to third-party conformity assessment; Annex III lists specific high-risk use cases. Assess the applicable requirements and the company’s provider or deployer role for each system.

United Kingdom (MHRA, NICE, etc.) – Post-Brexit, the UK's MHRA has pursued its own strategy. In April 2024, MHRA published an AI regulatory strategy with five pillars: safety/security, transparency/explainability, fairness/accountability, contestability/redress, and public trust. It explicitly frames **AI as a medical device (AIaMD)** if used for a medical purpose, meaning such software must meet the UK Medical Devices Regulations (an equivalent to EU MDR) (www.gov.uk). The MHRA signals that any AI/ML system for healthcare must be risk-classified and validated, but with **proportionate oversight** to encourage innovation (www.gov.uk) (www.gov.uk). As a regulator, MHRA is also adopting AI internally to improve efficiency (e.g. expediting reviews) (www.gov.uk) (www.gov.uk), which reinforces that it will expect well-understood evidence from regulated companies.

In June 2025, a UK government response to a Regulatory Horizons Council inquiry reaffirmed MHRA's commitment to a balanced approach, as detailed in the MHRA report "Impact of AI on Regulation of Medical Products" (www.gov.uk) (www.gov.uk). That report repeats: AI used for medical purposes must comply with existing frameworks (UDI, risk management, etc.) but notes that agile methods (e.g. iterative validation) may be advisable to fit the evolving nature of AI. Overall, UK guidance underscores stakeholder engagement, cross-industry dialogue, and readiness to adapt standards (e.g. updating ISO 13485 for SaMD) in light of AI's growth.

Other Jurisdictions – Canada (Health Canada) has collaborated with FDA on transparency principles. China's NMPA published [classification principles for AI-based medical software in July 2021](#). Device guidance must be distinguished from AI use in drug development; this overview does not establish an absence of drug-development guidance. Japan's PMDA has issued basic principles for AI/MDD (medical devices) and is developing guidance. Globally, organizations like WHO and OECD offer high-level AI ethics frameworks, but the operative rules are mainly national. Nevertheless, a trend toward international convergence is apparent: regulators increasingly issue joint statements (FDA–EMA, FDA–MHRA, etc.), Forum dialogues, and harmonized guidance (e.g. [ICH Q13 on continuous manufacturing of drug substances and drug products](#)).

Summary of Regulatory Trends: Across regions, several themes emerge. (www.fda.gov) (www.ema.europa.eu):

- **Risk-based approach:** AI systems are to be classified by risk (to patients, data integrity, etc.) and governance scaled accordingly (www.frontiersin.org) (www.fda.gov).
- **Transparency and Explainability:** Regulators insist on explaining AI decisions (at least to a reasonable degree), documenting algorithms' intended use, limitations, and providing human-interpretable information (www.fda.gov) (www.fda.gov).
- **Data Integrity and Privacy:** Applicable requirements continue to apply; assess jurisdiction, entity, intended use, and data or record type before assigning obligations. AI initiatives must maintain robust data governance (e.g. following ALCOA+ principles (www.auriacompliance.com)) and patient privacy at all times (talkingmedicines.com).
- **Lifecycle Oversight:** From development to post-market, AI needs continuous monitoring (e.g. software updates under change control, monitoring AI drift) as older static approvals evolve. FDA's device draft proposes a total-product-lifecycle approach [FDA device lifecycle draft](#).

- **Accountability and Ethics:** As a recommended governance control, companies should designate responsibility (possibly through AI ethics boards or officers) and build in human oversight (“human-in-the-loop”) to catch unforeseen failures ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).
- **Internal vs External Audit:** There is emphasis on internal ethics-based auditing (as seen at AstraZeneca) and readiness for external auditing/certification (especially under new laws) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)) (www.frontiersin.org). The **EMA–FDA principles** support a risk-based approach; **NIST AI RMF** supports multidisciplinary risk management.

These regulatory directives feed directly into enterprise AI governance requirements, as discussed in the sections below.

Core Governance Domains and Practices

Enterprise AI governance in pharma is multi-dimensional. Below we analyze key domains, each of which must be addressed through policies, processes, and controls that go beyond conventional IT or quality assurance.

1. Data Governance and Integrity

Definition & Standards: Data is the fuel of AI, and in pharma it is also a critical regulatory asset. For GxP data, apply integrity controls proportionate to intended use and data criticality; enterprise policy can extend these controls to other AI data. The **MHRA guidance** describes **ALCOA+**: data must be *Attributable, Legible, Contemporaneous, Original, Accurate*, and also *Complete, Consistent, Enduring, Available* (www.auriacompliance.com). If AI is used in GMP (Good Manufacturing Practices) processes, it must integrate with computerized system validation (CSV) to ensure data audit trails.

Privacy & Security: Patient health information often underlies AI models (e.g. EHR mining, trial recruitment). Assess privacy-law applicability for each deployment: **HIPAA covers covered entities and business associates**; assess GDPR and other laws by jurisdiction and data-processing context (talkingmedicines.com). This means strict controls on what data can be used (no unauthorized PHI), where models run (on-premises vs cloud, hashed data), and how logs/audit trails are kept. Generative AI (LLMs) raise new issues: even surrogate data (e.g. molecular structures) might be proprietary, so policies should specify the hosting, access, retention, and permitted data use of each deployment. A robust enterprise governance framework will define data classification (public, confidential, secret) and handling procedures, and restrict generation of AI content with regulated data (talkingmedicines.com).

Quality & Lineage: Beyond privacy, data used for AI training/testing must be validated for quality. In a regulated context, “garbage in, garbage out” is not acceptable. Governance requires a verifiable *data lineage*: every dataset should have provenance records, cleaning steps, and verification. This echoes traditional compliance: if trial data feed an AI, it must be the audited source data (SDTM, etc.), not unverified copies. A risk-based validation plan (per GAMP) should treat AI training as a “process” to qualify. Document data provenance and processing steps under the **EMA–FDA data-governance principle** ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

Controls & Recommendations:

- Classify and inventory all data assets used for AI.

- Enforce encryption, access logging, and retention rules as for any regulated data.
- Maintain versioning: if training data or models update, document changes in a change-management system.
- Evaluate access controls and privacy-enhancing methods for each use case.
- Train staff in data handling specifics for AI.

2. Model Development, Validation, and Quality Assurance

AI/ML systems are fundamentally software and must follow the SDLC (Software Development Life Cycle) with added ML specifics. Models must be **fit for purpose** and validated in context of intended use. ISPE GAMP 5 is industry good-practice guidance that can support risk-based computerized-system validation ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Indeed, researchers propose extending GAMP with specific steps for AI: e.g. formal model training plans, algorithm documentation, performance testing, and ongoing monitoring ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

[Appendix D11 of GAMP 5 Second Edition](#) addresses ML subsystem lifecycles. The [GAMP Artificial Intelligence Guide \(July 2025\)](#) is already published and is intended for use alongside GAMP 5 Second Edition. The following are proposed controls to tailor to the application and applicable requirements:

- Clear definition of the model *context of use* (what question it answers).
- Pre-specified performance metrics and acceptance criteria (analogous to assay validation).
- Use of hold-out test data or prospective validation to demonstrate model accuracy, sensitivity/specificity.
- Rigorous documentation: versioned code, annotated model cards (describing model architecture, training data, limitations) (www.frontiersin.org).

Pharma companies should treat AI model outputs like any quality attribute. For high-risk use (e.g. dosing recommendation, manufacturing control), their validation should be as stringent as a new assay or instrument. The FDA's draft guidance for AI in drugs underscores “model credibility”—trust that the model performs as intended (www.fda.gov). Therefore, sponsors must conduct comprehensive de-risking: adversarial testing, stress tests on edge cases, comparators to traditional methods. Any deficiencies must be documented and mitigated.

After deployment, periodic revalidation or monitoring is needed. AI models can “drift” as new data emerge. Governance should require tracking key performance indicators and triggers for re-training. For example, if a ML diagnostic tool's accuracy falls due to population changes, it must be flagged for review. [ICH Q14](#) concerns analytical procedure development and lifecycle management. Applying similar monitoring concepts to other AI systems is an adaptation, rather than a general ICH Q14 requirement for AI model monitoring or manufacturing process verification.

Controls & Recommendations:

- Integrate AI development into the quality system: require development plans, reviews, and QA sign-off stages.
- Use **checklists or standard operating procedures (SOPs)** for ML projects (proposed by practitioners ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/))).

- Employ tools like model cards and datasheets (proposed in ML literature) to standardize documentation (www.frontiersin.org).
- Perform thorough testing (unit tests, system tests, performance validation) and document results like any validated system.
- Have separate teams for model review (peer reviews by data scientists, code audits by IT security).
- Retain traceability: e.g. a regulated label file storing which model version was used to generate each result.

3. Risk Management

A risk-based approach underpins both regulation and governance. The importance of **risk categorization** is emphasized by experts: AstraZeneca, for instance, classifies AI systems as low/medium/high risk with proportionate controls (www.frontiersin.org). Low-risk might be an AI that formats label language (if no clinical judgment); high-risk could be one that predicts patient dosing. Governance policies should require upfront risk assessment for each use case, involving cross-functional stakeholders (quality, safety, IT, business) to determine the level of oversight needed.

During R&D and procurement, risk management means assessing vendor and technology. For AI tools acquired from external vendors or cloud platforms, due diligence on their governance (data handling, model validation) is needed. One approach is to treat AI procurement similar to critical software: require evidence of the vendor's own quality systems, or demand custom SLA terms for data/privacy. For each deployment, review hosting, access, retention, training reuse, vendor terms, and sensitive-data protections. Assess open-source and proprietary tools through the same procurement process instead of assigning privacy properties from an "open" or "closed" label. **NIST's Generative AI Profile** recommends policies for how third-party intellectual property and training data are used, stored, and protected. Governance should document the approved service and its specific data-use conditions.

Within development, risk management includes identifying failure modes (following approaches like Failure Mode and Effect Analysis). The **ICH Q9(R1)** framework on Quality Risk Management (2023) applies: identify hazards (e.g. biased output, cybersecurity threat), analyze their consequences, and implement mitigation plans (validation, bias audits, fallback processes).

Controls & Recommendations:

- Enforce a formal **AI risk assessment** step akin to hazard analysis, documented and approved by management.
- Use established risk frameworks: e.g. align with ICH Q9, ISO 14971 (for devices), or NIST AI RMF practices (www.fda.gov) (www.frontiersin.org).
- Scale controls: high-risk AI gets more oversight (e.g. human review of outputs, stricter change control, periodic revalidation).
- Monitor AI systems in operation (risk assurance) through audits or continual monitoring tools.
- Maintain a register of AI assets and their risk level, reviewed by QA or AI governance board periodically.

4. Transparency and Explainability

Pharma governance must ensure that AI decisions can be understood by regulators and, where relevant, by affected stakeholders (doctors, patients). Transparency is a recurring regulatory theme (www.fda.gov) (pmc.ncbi.nlm.nih.gov). For example, the FDA's Transparency Principles note that “logic” and “explainability” are aspects of transparency that may need to be communicated (www.fda.gov).

Practically, this means organizations should document how each AI model makes decisions in human-readable terms. While deep models are complex, explainable AI (XAI) techniques can provide surrogate explanations (feature importance, example-based reasoning) for critical cases. Pharma companies should invest in tools (like SHAP values, LIME) to generate audit letters explaining model outputs. Importantly, regulatory submissions should include information about model architecture, training data characteristics, and limitations. The FDA's draft guidance explicitly mentions providing “information needed for risk and benefit consideration” of AI models (www.fda.gov).

However, transparency also has logical limits. While “white-box” models are ideal, requiring full disclosure may conflict with intellectual property or security. Governance trade-offs are needed. As an enterprise recommendation, document approved systems, permitted inputs, and expected human review. These policies should make user responsibilities explicit.

Controls & Recommendations:

- Maintain detailed model documentation (“model cards”) that include inputs, outputs, training data summary, performance metrics, and known limitations (www.frontiersin.org).
- Implement logging for all automated decisions: record inputs, outputs, timestamps, and user interactions to create an audit trail.
- Use explainability tools especially for high-risk models, and include this in release notes.
- Communicate AI usage to end-users (e.g. workforce, clinicians): educate them on how AI outputs were generated and how to interpret confidence.
- Align with guidance: e.g. follow FDA transparency suggestions (www.fda.gov) and [ISO/IEC TR 24029-1:2021 on neural-network robustness assessment](#) as a separate robustness resource.

5. Ethical and Bias Considerations

In pharmaceuticals, fairness and ethics carry life-or-death weight. AI algorithms trained on biased or unrepresentative data can perpetuate health disparities or make unsafe recommendations. AI governance must explicitly guard against these risks. The World Health Organization has called for AI healthcare systems to promote **inclusivity and fairness**.

Governance policies should mandate bias evaluation: pre-deployment tests to detect disparate impact (e.g. by race, age, gender) and continuous monitoring for drift. For instance, an AI model for clinical triage should be audited to ensure it does not under- or over-treat any subgroup. These checks should be documented as part of the model validation reports.

Moreover, pharma companies should consider broader ethics. This includes avoiding misuse of AI (e.g. deepfakes for off-label marketing), respecting patient autonomy (AI suggestions should be advisory, not coercive), and ensuring accountability for decisions (clear escalation paths if AI error occurs). The [NIST AI](#)

RMF recommends incorporating diverse perspectives when identifying and managing AI risks.

Controls & Recommendations:

- Establish an ethics review or AI principles: e.g. AstraZeneca's bottom-up AI ethics principles which were collaboratively developed (www.frontiersin.org).
- Form cross-functional ethics committees (with ethicists, patient reps, etc.) to review high-impact AI projects.
- Require algorithmic impact assessments (AIAs), akin to data privacy DPIAs, for new AI deployments to identify ethical risks.
- Engage stakeholders (patients, HCPs) in the design and review of AI systems, particularly those affecting care.
- Provide clear channels for reporting AI-related concerns (whistleblowing) without fear of blame (www.frontiersin.org).

6. Organization and Culture

Technical controls alone are insufficient; governance is also an organizational challenge. Clear **ownership and accountability** are vital. Our sources note that lack of clear roles is a common barrier (pharmaconsulting.ai). Companies must decide who ultimately “owns” AI risks – often this spans IT, compliance, and the specific business function (e.g. R&D, manufacturing, sales) using the AI.

Best practice is to establish an enterprise AI governance body or council that brings together stakeholders from quality assurance, R&D, IT, privacy, regulatory affairs, and business leaders (www.fda.gov). Such a council can set company-wide policies (acceptable AI use cases, procurement rules, documentation standards) and also address cross-cutting issues (like a rapid response to a detected AI failure). The **NIST AI RMF** supports diverse teams and clear oversight roles. A governance council, ethics board, or designated officer is an implementation choice to tailor to the enterprise.

Leadership support is crucial. Studies show that **communication about AI governance is most effective when backed by senior executives** (www.frontiersin.org) (pmc.ncbi.nlm.nih.gov). This might involve an executive mandate on AI policy, or tying AI compliance to performance metrics for business units.

Organizational training is also key: employees from analysts to clinicians must learn how to use AI tools properly (knowing when to question an output). Without culture change and education, even the best policy fails.

Controls & Recommendations:

- Define and document roles: e.g. assign an “AI Compliance Officer” or expand duties of existing roles (e.g. Data Officer) to include AI oversight.
- Form an AI governance council or committee with cross-department representation.
- Develop clear policies on AI use cases: who can deploy generative models, what data can be used, etc. – and train employees on them.
- Encourage a culture of “responsible AI”: reward teams for proactive governance measures, make AI risk part of internal audits.

- Facilitate change management: when new AI policies launch, run workshops and incorporate feedback, as was done at AstraZeneca during its governance rollout (www.frontiersin.org).

7. Verification, Monitoring, and Audit

Even after deployment, AI systems require **continuous oversight**. This includes technical monitoring (for performance drift, security vulnerabilities) and process monitoring (audit compliance with governance procedures). For instance, if staff skip a mandatory AI decision review step, QA should detect and correct it.

Industry experts suggest formalizing audits of AI projects. The AstraZeneca case study (Frontiers, 2022) involved a third-party ethics audit, which took ~2,000 person-hours (www.frontiersin.org). Insights from that case show that *ethics-based auditing* faces the same governance hurdles as other audits: ensuring consistent standards, defining scope, and communicating results ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). Pharma companies should build audit clauses into AI procedures: e.g., internal audits could check whether model change control logs are complete, whether data used were documented, or whether outcomes were within expected bounds. For high-risk AI systems under the EU AI Act, the need for an independent conformity assessment body depends on the type of system, as explained in the [Commission's guidance](#). Assess the applicable route for each system.

Controls & Recommendations:

- Implement automated monitoring where possible: e.g. dashboards tracking model inputs/outputs, alerts for anomalous behavior, cyber threats.
- Periodically re-run validation tests or “red-team” challenges to confirm the AI still meets standards.
- Conduct regular internal audits of AI projects against a checklist (data governance, documentation, ethical sign-off).
- Consider external review to obtain feedback on AI governance. The AstraZeneca case study reports feedback on improving the company’s AI governance structures following external review (www.frontiersin.org).
- Update governance policies based on audit findings and evolving best practices.

8. Generative AI: An Emerging Frontier

In late 2022, generative AI (e.g. large language models like GPT) burst into mainstream consciousness. For pharma, gen-AI can draft reports, synthesize literature, and even suggest molecular structures. However, it introduces new governance wrinkles.

For generative AI, assess each approved deployment’s hosting, retention, access, training reuse, contractual terms, and sensitive-data protections. Do not infer these properties from whether model weights are open or proprietary. [NIST's Generative AI Profile](#) recommends an inventory of third parties with access to organizational content and approved technology/service-provider lists, together with policies for how third-party training data is used, stored, and protected.

Content creation also calls for verification: the NIST profile recommends deploying and documenting fact-checking techniques to verify generated information. Proposed enterprise controls include approved-tool lists, data-input restrictions, human review, and escalation of uncertain outputs.

Controls & Recommendations:

- Extend existing AI policy to explicitly cover generative tools. Specify which platforms are vetted and for what purposes.
- Forbid or block unapproved gen-AI use with controlled data. E.g., implement data leakage prevention software.
- If using gen-AI assistants, require people to annotate and archive all prompt-engineering and outputs (to maintain audit trail) (pharmaconsulting.ai).
- Include gen-AI risks in the same framework (bias, IP risk, misinformation). For complex outputs (e.g. predicting clinical signs), insist on high levels of validation and human oversight.
- Update training: educate staff that gen-AI can hallucinate or memorize, so use special caution in pharma contexts.

Comparative Table of Governance Domains

Table 2. Core AI governance domains in pharma, examples of concerns, and illustrative controls.

Domain	Key Considerations / Concerns	Example Controls & Policies (Illustrative)
Data Governance & Privacy	<i>Data integrity</i> (ALCOA+); proper data lineage; applicable privacy-law compliance; data quality	Data classification and access controls; encrypted storage; audit trails where required or adopted by enterprise policy; data-consistency checks; adversarial data cleaning. Use de-identified or synthetic data when possible (www.auriacompliance.com) (talkingmedicines.com).
Model Validation & Quality	<i>Model fit-for-purpose</i> ; reproducibility; software validation (GAMP 5); bias; versioning	Risk-based CSV : require model training documentation, code reviews, test plans. Use Good ML Practice guidelines (FDA) within GAMP workflows (pmc.ncbi.nlm.nih.gov) (pmc.ncbi.nlm.nih.gov). Maintain model history (versions, training logs). Performance benchmarks before release.
Risk Management	Formal risk assessment (context-of-use); classification (low/med/high); gap analysis	Conduct AI-specific risk assessments; integrate with ICH Q9 quality risk management. If assessed risk exceeds the organization's tolerance, define and document mitigating controls. Translate risks to SOP requirements (e.g. redundancy, fallback processes) (www.frontiersin.org) (www.fda.gov).
Transparency & Audit	IT/AI audit trails; explainability; documentation for regulators	Model cards/paper trail for each AI system (www.frontiersin.org). Maintain prompt logs and decision rationale. Provide users (HCPs, patients) with info on AI scope and limits (www.fda.gov) (www.fda.gov). Include AI usage in audit checklists.
Ethics & Fairness	Bias detection; equitable performance across demographics; ethical use cases	Regular bias/fairness audits (e.g. Demographic disparities tests). Establish AI ethics board to review high-impact systems. Adopt WHO/OECD principles to guide decision policies. Conduct privacy impact assessments for patient data.
Organizational Oversight	Roles & responsibilities; governance structures; training	Executive sponsorship (C-suite commitment). A proposed central AI governance council (cross-functional) or appointed AI officer , supported by clear oversight roles (NIST AI RMF) (www.fda.gov). Policies on approved tools, code of conduct. Ongoing staff training programs on AI use.

Domain	Key Considerations / Concerns	Example Controls & Policies (Illustrative)
Operations & Controls	Change management; monitoring during operation; incident response	Formal change control for updates to AI models/systems. Real-time monitoring dashboards; anomaly detection on outputs. Incident protocols: if AI error, revert to manual process, report to QA. Built-in kill-switches for unsafe outputs.
Legal & Compliance	Alignment with FDA/EMA guidance; contract management	Legal review of AI vendor contracts (IP, liability clauses). Compliance monitoring against evolving laws (NDA/reg changes). Assess applicable requirements; use relevant nonbinding guidance to inform regulatory-submission documentation.

Table 2: Major domains of AI governance in pharmaceutical enterprises, illustrating specific considerations and controls. Each control should be tailored to the company's risk profile and the AI application. Selected references are indicated.

Case Studies and Real-World Examples

Illustrating governance in action, we examine recent experiences from industry:

- AstraZeneca (Biopharma R&D)** – AstraZeneca's R&D Data Office has been a leader in operationalizing AI governance. An internal project (reported in *Frontiers in Comp. Sci.*, 2022) documented how AZ defined its AI scope and ethics principles by extensive bottom-up consultation. They classify AI projects by risk and integrate assessment into existing quality processes (www.frontiersin.org). For example, an NLP tool to draft a clinical report would undergo a moderate validation and include human review, whereas a low-risk use (like summarizing conference abstracts) needed lighter touch. AZ emphasized *education and incentives*: senior leaders championed AI governance, and teams were aligned with incentives to follow new policies (www.frontiersin.org). AZ also piloted an **ethics-based AI audit** (Q4 2021) with a third-party, devoting ~2,000 person-hours to it (www.frontiersin.org). Their lessons highlight that existing tools (model cards, impact assessments) must be coordinated in one overarching program (www.frontiersin.org). This case shows that significant upfront investment (time and people) is needed, but also that such governance can improve metrics like data security and brand reputation (www.frontiersin.org).
- Hypothetical communication-review workflow** – An illustrative pharmaceutical team could require a Regulatory Affairs specialist to review each AI-generated HCP letter before release. This is one proposed risk-control step within a separately assessed promotional and privacy process; it does not establish a measured company outcome or guarantee compliance.
- AI evidence in drug submissions** – FDA's [January 2025 draft guidance](#) describes a risk-based framework for assessing AI model credibility for a particular context of use. Sponsors should distinguish these nonbinding draft recommendations from binding legal requirements.

These examples illustrate proposed controls and the need to tailor evidence and oversight to the application.

Implications and Future Directions

Looking ahead, AI governance in pharma will continue to evolve. Several trends should be anticipated:

- **Implementation of Enacted Law:** Prepare an applicability assessment against the [EU AI Act's legal classification criteria](#) and the [Commission's phased timeline](#). Assess the relevant provider/deployer role and conformity-assessment route; do not use an unsupported percentage of development cost as a current certification budget benchmark. The 2023 U.S. Executive Order 14110 was [rescinded on 20 January 2025](#) and does not establish a current FDA “AI Act-like” framework.
- **Global Harmonization:** There is momentum for international standards. On one hand, differing laws (EU vs U.S. vs China) could complicate global programs. On the other, initiatives like FDA-EMA collaboration signal a push toward harmonization. Industry groups (EFPIA, BIO, IFPMA) are advocating for flexible, innovation-friendly rules that nonetheless ensure patient safety (www.ema.europa.eu). Companies should engage in these dialogues to shape regulations earnestly.
- **Generative AI Governance:** Review the specific laws and contractual terms applicable to each generative AI deployment, and document controls for scientific accuracy, sensitive data, and human oversight. Treat proposed rules and enterprise recommendations separately from current legal duties.
- **Technology Solutions for Governance:** Just as blockchain emerged to solve traceability, we may see new tools. For instance, AI “audit algorithms” could monitor other AI; watermarking for provenance of data; standardized model registries. NIST and others are also working on automation-friendly guidelines. Companies should watch for frameworks like **ISO/IEC 42001** (governance of AI) and alignment with federal Trustworthy AI strategies (www.nist.gov).
- **Workforce and Culture:** As AI diffuses, more non-technical staff (clinicians, quality personnel, marketers) will use AI tools. Governance cannot rely solely on specialists. Enterprises must invest in educating the workforce on AI risks and best practices. Some propose that future pharma professionals may require certification in AI governance, much like cGMP certification today.
- **Sustainability and Ethics:** Broader questions will enter governance debates. For example, what about environmental impact of large-scale computing, or equity in global health (AI models trained only on Western data)? Regulatory focus might expand to these issues, especially as ESG concerns rise.
- **Business Value vs Compliance:** Finally, effective AI governance will be seen not just as a compliance burden but a value driver. As noted in AstraZeneca’s case, robust AI governance can improve data quality, speed approvals, and even brand image (www.frontiersin.org). Early adopters who integrate governance smartly may gain competitive advantage by accelerating safe AI adoption and avoiding costly recalls or fines.

Conclusion

Enterprise AI governance for regulated pharma is complex but imperative. Unlike consumer tech, missteps here can directly endanger patients and violate law. This report has shown that addressing AI governance demands an **integrated approach**: technical controls, risk management, policy, and culture must all work in

concert. Recognizing this, leading regulators and industry bodies now provide detailed guidance: from FDA's drug AI framework to the EMA–FDA joint principles. Companies must assimilate these into actionable programs.

Core principles for pharma AI governance include: maintaining data integrity and privacy at the highest standard (www.auriacompliance.com) (talkingmedicines.com); employing risk-based validation and transparency (www.fda.gov) (www.fda.gov); and building organizational structures for accountability (www.frontiersin.org). Operationally, this means extending GxP-quality processes (change control, audits) to AI, and treating AI outcomes with the same rigor as traditional data. We provided tables summarizing these domains and how controls might look in practice.

The historical context shows steady progress: from rudimentary AI pilots to mature pipelines. Today's recommendations are not merely theoretical – they reflect the reality that 2026 and beyond will see a proliferation of AI in regulated functions. Companies that proactively embed governance will reap not only compliance, but trust and efficiency gains. For example, ensuring AI approvals for a candidate drug might be faster if model validation and documentation are already in place. Conversely, ignoring governance invites derailment of AI projects or worse – patient harm.

As we look to the future, one thing is clear: enterprise AI governance in pharma is no longer optional. It is an essential component of good business and scientific practice. By adhering to robust standards and evaluating experience within the intended context of use, pharmaceutical companies can leverage AI's promise safely. This report, grounded in extensive research and official guidance, aims to serve as a detailed roadmap for organizations navigating this new frontier. Recommendations and hypothetical examples are proposed enterprise practices; they require application-specific assessment.

References: Sources are linked inline in the relevant passages and tables. Regulatory documents, voluntary frameworks, industry estimates, and author recommendations have different evidentiary and legal roles. For organizational oversight, [NIST AI RMF](#) supports multidisciplinary perspectives and explicit oversight responsibilities. No numbered footnote bibliography is used.

Source links

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

1. <https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2022.1068361/full>
2. <https://www.auriacompliance.com/gmp-blog/data-integrity-and-ai-integration-key-considerations-for-compliance-in-gmp-pharmaceutical-manufacturing>
3. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9152664/>
4. <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions>
5. <https://www.ema.europa.eu/en/about-us/how-we-work/data-regulation-big-data-other-sources/artificial-intelligence>
6. <https://www.fda.gov/medical-devices/medical-devices-news-and-events/cdrh-issues-guiding-principles-transparency-machine-learning-enabled-medical-devices>
7. <https://www.fda.gov/medical-devices/software-medical-device-samd/transparency-machine-learning-enabled-medical-devices-guiding-principles>

8. <https://www.ema.europa.eu/en/news/reflection-paper-use-artificial-intelligence-lifecycle-medicines>
9. <https://www.gov.uk/government/news/mhras-ai-regulatory-strategy-ensures-patient-safety-and-industry-innovation-into-2030>
10. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7892696/>
11. <https://www.ema.europa.eu/en/scientific-guidelines/use-artificial-intelligence-ai-medicinal-product-lifecycle>
12. https://www.ema.europa.eu/system/files/documents/other/2026-01-13_ai-guiding-principles_ema-fda_en.pdf
13. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
14. <https://ispe.org/pharmaceutical-engineering/july-august-2022/ai-governance-and-qa-framework-ai-governance-process>
15. <https://www.fda.gov/science-research/science-and-research-special-topics/artificial-intelligence-and-machine-learning-aiml-drug-development>
16. <https://pharmaconsulting.ai/ai-governance-pharmaceutical-industry/>
17. <https://www.mckinsey.com/industries/life-sciences/our-insights/generative-ai-in-the-pharmaceutical-industry-moving-from-hype-to-reality>
18. https://assets.publishing.service.gov.uk/media/5aa2b9ede5274a3e391e37f3/MHRA_GxP_data_integrity_guide_March_edited_Final.pdf
19. <https://www.hhs.gov/hipaa/for-professionals/covered-entities/index.html>
20. <https://talkingmedicines.com/2024/12/top-8-governance-considerations-for-ai-in-pharma/>
21. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/part-11-electronic-records-electronic-signatures-scope-and-application>
22. <https://www.nist.gov/artificial-intelligence/executive-order-safe-secure-and-trustworthy-artificial-intelligence>
23. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
24. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>
25. <https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>
26. <https://ispe.org/pharmaceutical-engineering/updated-gamp-gpg-incorporates-ai-and-open-source-software>
27. <https://ispe.org/publications/guidance-documents/gamp-guide-artificial-intelligence>
28. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing>
29. <https://healthit.gov/regulations/hti-rules/hti-1-final-rule/>
30. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological>
31. <https://www.statnews.com/2025/01/06/artificial-intelligence-use-drug-development-fda-draft-guidance/>
32. <https://www.gov.uk/government/publications/impact-of-ai-on-the-regulation-of-medical-products/impact-of-ai-on-the-regulation-of-medical-products>
33. https://english.nmpa.gov.cn/2021-07/08/c_660267.htm
34. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/q13-continuous-manufacturing-drug-substances-and-drug-products>
35. <https://www.fda.gov/media/161202/download>
36. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
37. <https://www.iso.org/standard/77609.html>
38. <https://www.ema.europa.eu/en/news/ema-fda-set-common-principles-ai-medicine-development-0>

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.