



INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Pharma AI Vendor Due Diligence: A Validation Checklist

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 · 2026-02-14 · pharma ai · vendor due diligence · ai validation · gxp compliance · fda regulations · life sciences · machine learning · data security · risk assessment

Original article: <https://intuitionlabs.ai/articles/pharma-ai-vendor-due-diligence-checklist>



In this article

1. Executive Summary
 2. Introduction and Background
 3. The AI Hype Cycle and Pharma Context
 4. Core Areas of Due Diligence
 5. Background: Historical and Regulatory Context
 6. Technical Due Diligence
 7. Case Studies and Illustrative Examples
 8. Legal, Compliance, and Contractual Due Diligence
 9. Security and Privacy Due Diligence
 10. Performance and Integration Considerations
 11. Implications and Future Directions
 12. Conclusion
 13. References
-

Executive Summary

As artificial intelligence (AI) rapidly transforms pharmaceutical R&D, manufacturing, and commercialization, life science companies are increasingly turning to AI vendors to gain competitive advantage. Yet the promise of AI often collides with reality: vendor marketing ploys, unverified performance claims, and complex regulatory constraints frequently lead to unmet expectations or outright failures. Consequently, pharmaceutical buyers must conduct **rigorous due diligence** to validate AI vendor claims before signing contracts or deploying solutions. This report delivers a comprehensive checklist and framework for that due diligence, tailored to the unique requirements of the pharmaceutical industry.

Key findings and recommendations include:

- **Heightened Stakeholder Scrutiny:** AI is now an “immediate priority” for most Big Pharma (85% see AI as urgent (www.fiercepharma.com)), yet analysts warn of “overheated marketing” and hype surrounding AI drug discovery (www.statnews.com). Buyers must reconcile optimism with skepticism by demanding evidence for all vendor claims and by involving multidisciplinary stakeholders (R&D scientists, regulatory, IT, etc.) in evaluations (www.clinicalleader.com) (www.pharmoutsourcing.com).
- **Data & Model Transparency:** Vendors must disclose core details about data sources and model validation. Do not assume that a pharmaceutical or clinical specialization establishes superior performance. A [2026 Nature Medicine comparison](#) found frontier general-purpose LLMs outperformed the tested clinical AI tools in three evaluations. Treat this as counterevidence to blanket specialist superiority, rather than a universal ranking. Buyers should require blinded comparisons against both general-purpose and specialist alternatives on representative tasks, using the proposed model versions, data, and deployment configuration. Insist on documentation of training, versioning, and performance metrics on representative data.

- **Regulatory Compliance:** AI tools in pharma often fall under regulations (e.g. FDA, EMA, GxP, HIPAA). Buyers should verify that vendors follow relevant guidelines: e.g. [FDA's AI/ML draft frameworks](#), GAMP 5, and ISO standards for medical devices and software. Audit trails, validation documents, and risk assessments are critical (www.pharmaceuticalonline.com) (pmc.ncbi.nlm.nih.gov).
- **Performance Verification:** Therapeutic decisions and manufacturing processes demand high accuracy. Vendors claiming high prediction accuracy must provide independent test results (not cherry-picked metrics). Performance metrics should be defined in advance and validated in multiple scenarios (pmc.ncbi.nlm.nih.gov).
- **Bias and Ethics:** Pharma must ensure patient safety and equity. AI models trained on biased clinical data can produce discriminatory outcomes. Buyers should query vendors about bias testing and mitigation, and require transparency around datasets.
- **Security & IP Protection:** Generative AI vendors often encourage customers to input confidential data. Contracts must explicitly restrict vendor reuse of sensitive prompts (pharmaphorum.com). Buyers should check encryption standards (e.g. HIPAA compliance, [ISO 27001](#)), incident response plans, and segregated data environments (pharmaphorum.com) (www.clinicalleader.com).
- **Contractual Safeguards:** Vendors typically aim to limit liability. Pharma buyers should negotiate strong indemnities, performance-based Service Level Agreements (SLAs), and clear IP ownership. They should demand vendors maintain professional liability insurance and agree to rapid correction of failures (innovaccor.com) (pharmaphorum.com).
- **Implementation and Change Management:** Even the best AI can fail without user buy-in. Case studies show that tools without alignment to scientists' workflows go unused (www.pharmoutsourcing.com) (www.pharmoutsourcing.com). Due diligence should include pilot studies and UX evaluations to ensure tools meet real needs.

This report develops these findings into a structured **due diligence checklist**, supported by academic and industry sources. We analyze historical and current examples, discuss regulatory/risk contexts, and outline future trends. Tables summarize key questions buyers should ask. In sum, responsible AI procurement in pharma demands depth: verifying technical claims, ensuring compliance, and safeguarding patient and corporate interests.

Introduction and Background

The AI Imperative in Pharma

By 2026, AI has become ubiquitous in life sciences. Every major pharmaceutical firm is exploring AI for drug discovery, clinical trials, manufacturing, and even marketing. One survey found 70–85% of top pharma executives consider AI an “immediate priority” (www.fiercepharma.com) and are increasing budgets for AI projects. Define Ventures reported that **85%** of Big Pharma leaders view AI as urgent, with over 80% raising AI spending (www.fiercepharma.com). Use cases range from *in silico* drug screening to [automating regulatory paperwork](#) and mining [real-world evidence](#). Analysts estimate that scaled AI could generate **\$254 billion** in extra annual profits for pharma by 2030 (www.baringa.com).

However, with massive hype comes risk. Stat News highlights an “unending stream of overheated marketing” in AI drug discovery, stressing that vendors often promise breakthroughs that may arrive years later or never (www.statnews.com). A 2024 industry survey found 83% of pharma professionals called AI “overrated,” despite widespread deployment (www.fiercepharma.com). Real-world AI implementation lags – a Gartner study shows while 92% of life science companies have pilots, only 30% have scaled more than six projects (www.baringa.com). Many initiatives stall due to data issues or governance gaps, as analysts repeatedly note (www.baringa.com) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)).

Standard due diligence processes (e.g. vendor questionnaires, tech evaluations) often fail to capture AI-specific risks. AI-enabled tools are not like typical software; they may rely on opaque models trained on proprietary data, produce unpredictable outputs (hallucinations), or degrade with data drift (www.pharmaceuticalonline.com) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). Pharma’s regulatory environment further complicates matters. For computerized equipment used in drug manufacturing, **21 CFR 211.68** bases the degree and frequency of input/output verification on system complexity and reliability. Identify each tool’s regulated activity and applicable records before setting validation and change-control duties; avoid treating every pharmaceutical AI system as subject to an identical GMP mandate.

In this climate, a pharma-buying organization must “**trust but verify**” every claim an AI vendor makes (www.pharmaceuticalonline.com). This report provides the historical context, current best practices, and future outlook for that verification process. We draw on peer-reviewed studies, regulatory guidelines, industry reports, and expert commentary to assemble a detailed AI due-diligence checklist specifically for pharmaceutical applications.

Defining AI Vendor Claims and Due Diligence

AI vendors may claim a wide array of capabilities: e.g., “90% accuracy in safety signal detection,” “FDA-compliant regulatory review,” “natural language experts for clinical documentation,” or “halving drug screening time.” Each claim carries assumptions about data, algorithms, and context. Due diligence is the systematic process of testing those assumptions and gathering evidence to substantiate (or refute) the claims. It spans multiple domains:

- **Technical validation** of algorithms and data (model performance, robustness, data provenance).
- **Regulatory and compliance assessment** (alignment with FDA/EMA guidance, data privacy laws).
- **Security assessment** (cybersecurity posture, data protection mechanisms).
- **Business and legal review** (vendor’s track record, contractual terms, liability).

In regulated pharma, due diligence is not optional. Unvalidated AI can result not only in business loss but also patient harm, regulatory violations, intellectual property leakage, or litigation. A rigorous AI vendor due diligence process protects all these interests.

In the following sections, we dissect each dimension of due diligence. We include multiple perspectives (legal, technical, business, and user) and illustrate with examples. Citations are provided for all substantive points. For clinicians or managers new to AI, we begin with a primer on the stakes and recent learnings, before diving into specific checklist items.

The AI Hype Cycle and Pharma Context

Evaluating Vendor Claims

Evaluate vendor marketing against the proposed intended use and independently reviewed evidence. Promotional claims should prompt questions about the test design, comparison group, and limitations.

- **Generic vs Domain-Specific AI:** Model specialization alone is insufficient evidence of suitability. In a [2026 Nature Medicine study](#), frontier LLMs outperformed the tested clinical AI tools across medical knowledge, clinician alignment, and real clinical queries. This is counterevidence to categorical specialist superiority, not proof that general-purpose models always win. Require blinded, representative comparisons of both alternatives for each task, model version, data source, and deployment configuration; define acceptance criteria before testing.
- **Over-Promising Results:** Pharma leaders report that many AI vendor success stories lack supporting data. The CEO of Insitro (a prominent AI biotech) warns that people “*think [AI] is going to happen tomorrow*” (www.statnews.com), yet true breakthroughs remain years away. STAT’s Casey Ross and Matthew Herper observe an “unending stream of overheated marketing” in AI drug discovery (www.statnews.com). This echoes the broader tech phenomenon where startups tout “AI cures X in 50% less time” without peer-reviewed evidence.

The takeaway: buyers must adopt a critical stance. Any claim about accuracy, speed, or ROI should be verified against data and aligned with expert judgment. Marketing materials should be a starting point for investigation, not an endpoint of trust.

Industry Adoption Trends and Skepticism

Despite skepticism, pharma’s investment in AI has grown sharply. A 2025 survey of global pharma executives found **70–85%** of respondents deem AI an “immediate priority,” with most expanding AI budgets (www.fiercepharma.com). Companies are now pairing external vendors with internal initiatives; 40% combine in-house and vendor solutions (www.fiercepharma.com).

However, the same survey notes that investment is targeted at “low-risk, high-efficacy” areas (like medical writing) initially (www.fiercepharma.com). A separate [Klick Health/Momentum Events promotional-review survey](#) collected 40 online responses from U.S. pharmaceutical marketing/promotional-review professionals between July 11 and October 8, 2025; 30% identified Regulatory as their primary review role. In that sample, 65% did not trust AI to create regulatory compliance submissions, while 40% cited hallucinations and 20% cited lack of traceability/audit trails. These findings describe this small promotional-review sample and should not be generalized to all pharmaceutical compliance or drug-submission professionals.

In practice, many companies create AI governance committees to oversee ethics and data policies. Define Ventures reported that ~80% of firms have dedicated AI governance structures in place (www.fiercepharma.com). This governance environment underlines why due diligence is central: vendor decisions carry not just technical risk, but also ethical and compliance risk under these new oversight regimes.

Core Areas of Due Diligence

To systematically evaluate an AI vendor, pharmaceutical purchasers should group questions into major domains. Below is a high-level outline; detailed criteria appear in subsequent sections and the accompanying tables.

1. Technical Architecture and Model Validation.

- *Questions:* How is the AI/ML model constructed? Is it proprietary, open-source, or a hybrid? What algorithms and features are used? Has the model been validated on independent test sets that resemble our data?
- *Rationale:* Transparent understanding of the model helps assess faith in claims. For example, if a vendor touts *95% accuracy*, know whether that is on training data, test data, or unseen clinical records. Validation must match intended use to avoid inflated metrics ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)).

2. Data Provenance and Quality.

- *Questions:* What data was used to train and test the model? Is the vendor's data representative of our patient populations and therapeutic areas? How was data labeled and cleaned?
- *Rationale:* "Data is the foundation of AI" (www.pharmaceuticalonline.com). Flawed or biased training data means flawed predictions. PHARMA buyers must ensure training sets meet ALCOA+ standards (Attributable, Legible, Contemporaneous, Original, Accurate, etc) (www.pharmaceuticalonline.com) and correspond to the use-case population (www.pharmaceuticalonline.com).

3. Performance Metrics and Testing.

- *Questions:* What metrics does the vendor report (e.g. accuracy, AUC, precision/recall)? Are these metrics sensitive to class imbalance (common in safety data)? What are the false positive/negative rates?
- *Rationale:* Standard accuracy can be deceptive. In safety monitoring, for instance, serious events may be only 10% of cases (www.vitrana.com). Vendors should provide balanced metrics and explain them in context ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)).

4. Regulatory Compliance and Validation.

- *Questions:* Does the AI solution fall under medical device or software regulations (e.g. FDA's SaMD guidelines, EU Medical Device Regulation)? Has the vendor performed a Computerized System Validation or equivalent?
- *Rationale:* In regulated pharma, digital systems must be validated to show they are "fit for intended use" (www.pharmaceuticalonline.com). For example, any AI used in pharmacovigilance must comply with GxP and CSV guidelines ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). Proper validation documentation is as important as technical performance.

5. Security, Privacy, and Data Usage.

- *Questions:* How does the vendor protect data? Are they HIPAA/GDPR compliant? What contractual rights do they claim to customer data and outputs (often a contentious issue with GenAI)?
- *Rationale:* AI systems often process sensitive IP (e.g. patient data, proprietary formulas). Contracts must explicitly restrict vendors from storing or reusing user data for other purposes (pharmaphorum.com). Security certifications (SOC2, ISO 27001) and incident response plans are key evidence of maturity.

6. Ethical Bias and Explainability.

- *Questions:* Has the vendor tested for algorithmic bias across demographics (age, gender, ethnicity)? Does the AI provide explanations for decisions (e.g. feature importance, source references)?
- *Rationale:* Biased predictions can harm patient safety and violate regulations. Buyers should require descriptions of bias-mitigation strategies and evidence of explainability (www.fiercepharma.com) (discover-pharma.com).

7. Business and Operational Fit.

- *Questions:* How will the AI integrate into existing workflows (e.g. LIMS, SAP, EHR)? What changes or training will be needed? Does the vendor have references or prior deployments in pharma?
- *Rationale:* Even a high-performing model is useless if scientists won't use it (www.pharmoutsourcing.com) (www.pharmoutsourcing.com). A successful solution demands good UX design and alignment with user jobs. Also, as one due diligence framework notes, the ROI of an AI solution must be assessed in context of enterprise goals (pmc.ncbi.nlm.nih.gov).

8. Contractual and Legal Terms.

- *Questions:* What warranties and liabilities does the vendor accept? Are service levels tied to performance (uptime, throughput)? Who owns the IP of model improvements? What happens if the vendor goes bankrupt or changes strategy?
- *Rationale:* AI vendors often try to limit liability for “model outputs”. Pharma buyers must negotiate “skin in the game” – e.g. indemnities, insurance, and termination rights (pharmaphorum.com) (www.clinicalleader.com). Crucially, due diligence includes legal review to shift undue risk.

In the sections that follow, we delve into each pillar with detail and supporting evidence. We discuss the historical context of why each question matters (with examples), cite relevant guidelines and case studies, and suggest practical verification steps. Two summary tables at the end condense these considerations into actionable checklists.

Background: Historical and Regulatory Context

Before diving into specifics, it is useful to review the evolving context of AI in pharma that shapes due diligence requirements.

Regulatory Landscape

Pharmaceutical AI can implicate different regulatory regimes. Begin with a documented map of the intended use, jurisdiction, regulated activity, records, and responsible parties; do not infer legal status from the vendor's AI label.

- FDA (USA):** Separate a medical-device intended use from AI supporting drug development or labeling documents. For a device, request the actual regulator decision, authorized indications, product/version, and limitations. Verify the pathway accurately: **510(k) clearance** and **De Novo classification and marketing authorization** are distinct; neither establishes approval of drug labeling or a drafting service. For drug or biologic regulatory support, review the sponsor's applicable submission evidence and the context of use. FDA's **January 2025 drug/biologic AI guidance** remains draft guidance, not for implementation, with nonbinding recommendations; its framework addresses model credibility for a particular context of use. The completed **Software Pre-Cert pilot** should not be treated as a current vendor-certification route.
- EU AI Act:** This is an enacted law with phased applicability, rather than an upcoming proposal. Classification depends on intended use. Under **Article 6**, the product route requires both an Annex I product or safety component and the applicable third-party conformity-assessment condition; Annex III lists other high-risk uses, subject to its specified exceptions. Pharmaceutical analytics are not automatically high-risk. Document the analysis, including any claimed **research/development exclusion under Article 2**; do not extend that exclusion to routine deployment. Consult the Commission's **current application timeline** for the provisions relevant to the product and use.
- GxP, GAMP, and electronic records:** Map the system to applicable regulated activities and predicate-rule records before setting validation duties. Treat GAMP as industry guidance, rather than a binding regulation. FDA's **Part 11 scope and application guidance** distinguishes applicable electronic records from other records and recommends a documented, risk-based validation approach. Identify the enforceable predicate-rule requirements separately from guidance recommendations, and agree change controls, testing, and record-integrity evidence appropriate to that scope.
- Data Privacy:** HIPAA security risk analysis and GDPR DPIAs are different assessments. For HIPAA-covered entities and business associates, assess risks to electronic protected health information and document the analysis, using **HHS guidance on risk analysis**. Under GDPR, a **DPIA is required where processing is likely to pose high risk to individuals' rights and freedoms**; novelty alone is not the test. Require a documented lawful basis, permissions, safeguards, and contracts for each dataset and use. Avoid treating consent plus de-identification as a universal combined requirement; examine the applicable **HIPAA research provisions** when research PHI is involved.
- Liability and contracting:** Ask counsel to identify the applicable product-liability, privacy, and other legal duties for the jurisdiction and use. Negotiate responsibilities, indemnities, insurance, and escalation procedures against those duties rather than relying on an unspecified proposed AI instrument.

For every vendor claim of authorization or compliance, request the underlying decision or documented legal assessment. Confirm that the exact product, intended use, records, and deployment are within scope before accepting the claim.

Complexity of AI-Pharma Sector

Investors and analysts note that pharma AI is exceptionally complex. One industry article bluntly states: “*AI-Pharma companies are 100 times as complex as FinTech companies. Methodologies used to assess them should be 100 times as rigorous.*” ([healthmanagement.org](https://www.healthmanagement.org)). Why? Developing AI for drug discovery integrates biology, chemistry, regulatory science, and cutting-edge machine learning. Senior pharma technologists often describe AI startups as mysterious “black boxes” to outside investors ([healthmanagement.org](https://www.healthmanagement.org)). Even venture capitalists have had to devise new due diligence frameworks for AI-driven biotech, assessing hundreds of parameters describing technical maturity, team expertise, data assets and regulatory strategy ([healthmanagement.org](https://www.healthmanagement.org)).

For the buyer (who is not selling but procuring AI systems), similar scrutiny is warranted. Unlike purchasing a chemical reagent, deploying an AI model means betting on intangible algorithms and data pipelines. Poor diligence can easily lead to a solution that “**technically works but doesn’t meet user needs**” as one expert warns (www.pharmoutsourcing.com) (www.pharmoutsourcing.com).

This guide takes that advice to heart: we advocate deep, cross-disciplinary review of AI vendors in pharma. It is only by understanding the scientific and regulatory intricacies, alongside technical performance, that a buyer can legitimately trust an AI product.

Technical Due Diligence

This section addresses core technical claims: data, model, and performance. We advise on how to evaluate the **algorithmic validity** of an AI vendor’s offerings.

Data Quality and Scope

Key Questions: What datasets underlie the AI tool? Have they been independently audited for accuracy? Are they representative of the intended patient population, therapeutic areas, or laboratory environment? Were the data collected ethically and in compliance with regulations?

Why It Matters: In AI, the adage “garbage in, garbage out” is literally true. Pharmaceutical data often present challenges: small sample sizes, noisy lab records, heterogeneous sources (electronic health records, genomic assays, real-world evidence, etc.). Generic data cleaning tools may not suffice. Huysentruyt et al. identify potential biases, completeness, and data sources as data-quality risks in AI-system validation ([pharmacovigilance validation framework](https://www.pharmacovigilancevalidationframework.org)).

Regulators implicitly demand high data integrity. The ALCOA+ framework, borrowed from traditional pharma practices, applies to training data: every datum should be attributable, accurate, original, etc. (www.pharmaceuticalonline.com). Buyers should insist vendors document data lineage (where each record came from) and handling (how it was normalized, anonymized, etc.).

In practice, ask vendors for:

- **Data Descriptions:** summaries of data sources, volumes, and types. E.g. if an AI monitors lab machines, which factories? If for text completion, what clinical notes or lab standards (CDISC, HL7) were used?
- **Data Audits:** reports on data quality control. Have they run sanity checks or bias audits (e.g. demographics of patient data)?
- **Coverage:** evidence that critical sub-populations or rare conditions aren't missing. For example, if an AI claims to classify side-effects, it should have been trained on ADE reports spanning the relevant geographies and patient demographics, or supplemented by synthetic cases if needed.

Documented assurance of data integrity aligns with GxP: systems must have traceable input/output records. If possible, the buyer may request a *Data Protection Impact Assessment (DPIA)* or similar from the vendor that clarifies legal bases for each dataset.

Model Architecture and Validation

Key Questions: How is the AI model built and validated? Are the underlying algorithms explained? Have external experts reviewed or replicated the results? Is there version control and retraining governance for the model?

Why It Matters: Models in pharma AI range from simple statistical regressions to deep neural networks or ensemble methods. The choice of technique impacts transparency and robustness. A heuristic is: *“more complex ≠ always better”*. For example, a random forest might suffice for batch process optimization, whereas drug imaging may require convolutional nets.

Buyers should inquire:

- **Model Type:** Ask if it's supervised learning, unsupervised, reinforcement, etc. If it is a neural net, is the architecture described (e.g. layers, hyperparameters)?
- **Training Process:** How was the model trained? Was the training hyperparameter search automated or manual? What computing resources (GPU clusters, distributed training) were used?
- **Validation Data Split:** Very importantly, were training/validation/test sets strictly separated? If robust machine learning practices are followed, there should exist held-out test sets never seen during development, preventing “data leakage” (www.pharmaceuticalonline.com) (pmc.ncbi.nlm.nih.gov). The vendor should commit that any performance metrics are from fresh test data.
- **Reproducibility:** Ideally, the buyer should have confidence that the results are reproducible. That could mean obtaining a technical appendix or whitepaper describing the algorithm, or even requesting a joint workshop where vendor data scientists demonstrate training on a subset of data.

If the vendor has published results (e.g. in a journal or conference), review those publications. Independent benchmarking (e.g. via Kaggle challenges or known datasets) can be telling: did the model truly beat existing baselines on problems similar to yours?

(For example, many so-called “AI” drug candidates are never peer-reviewed; investors and buyers should be skeptical of unfounded claims until results appear in trusted journals or are validated in independent trials.)

Performance Metrics

Evaluation of an AI model's performance must go beyond headline accuracy. Key due diligence steps include:

- **Define Appropriate Metrics:** Ensure that metrics match the use case. For classification tasks with class imbalance (common in safety signal detection, where adverse events are rare), accuracy alone can be misleading (www.vitrana.com). Prefer metrics like Area Under ROC (AUC), F1-score, or recall at a fixed precision. For predictive models in medicine, sensitivity (catching true positives) may be more critical than overall accuracy. Vendors should justify why chosen metrics reflect business needs, and provide confusion matrices if possible.
- **Benchmarking and Baselines:** Insist on comparisons to baseline methods. If a vendor claims “AI algorithm identifies compounds 50% faster,” they should compare against conventional methods (statistical models, or manual lab screening). Similarly, if the vendor touts error rates, have they been compared to human performance or other algorithms?
- **Stress Testing:** Robustness tests, such as running the model on out-of-distribution data or introducing noise, can reveal reliability. Buyers might request a summary of such tests. Regular performance monitoring (e.g., tracking AUC drift over time) is also key, as data and real-world conditions change (www.pharmaceuticalonline.com).

Citing specifics: the pharmacovigilance literature and FDA guidances stress that “*performance metrics must be proactively defined given the system’s intended use*” (pmc.ncbi.nlm.nih.gov). For high-risk tasks (diagnostics, patient dosing), vendors should specify error thresholds (e.g. maximum false negatives allowed) upfront (www.pharmaceuticalonline.com) and document how performance was measured.

Explainability and Transparency

Key Questions: Are model decisions explainable? If the model is a black box, what tools (like SHAP or LIME) does the vendor use to interpret outputs? Does the vendor provide documented reasoning or source references for AI-generated predictions?

Why It Matters: Pharma is highly risk-averse. Regulators and end users demand some level of rationale for AI decisions. For instance, if an AI suggests a drug motif, chemists want to know *why* (which structural features it “found” relevant). In clinical AI, physicians expect at least confidence scores or reference cases. Lack of explainability can kill trust: Buyers should test whether users can understand and verify the rationale behind outputs in the intended workflow.

Due diligence tip: Ask for examples of explainability reports. For sample cases, does the AI output include tracebacks (e.g. source text snippets, or feature importances)? Verify consistency: if given the same input twice, does the model produce stable explanations? Some vendors claim to avoid “hallucinations” by design. Probe what that means technically.

(Example: A healthcare AI platform might highlight the sentences of a medical text that triggered an AI diagnosis. The buyer should test if those highlights actually match clinical reasoning, rather than being random artifacts.)

Security and Data Protection

Key Questions: How does the vendor secure data in transit and at rest? Are there third-party security certifications? Does the vendor allow on-premises deployment or only cloud-based? How are encryption keys managed?

Why It Matters: Pharmaceutical data (clinical trials, molecular structures, patient info) are highly sensitive. Security breaches risk not just IP loss but regulatory fines and patient harm. AI vendors must demonstrate enterprise-level security:

- **Standards Compliance:** Look for ISO/IEC 27001, SOC 2 Type II, or FedRAMP certifications. These indicate formal security processes.
- **Infrastructure Controls:** Does the solution run on public cloud, private cloud, or on-premises? If cloud, which provider (AWS, Azure, etc) and region? For global operations, ensure data residency rules are respected.
- **Data Segregation:** The vendor should explicitly segregate each customer's data. Shared multi-tenancy must not allow data leakage.
- **Access Management:** Role-based access control, multi-factor authentication, and strict audit trails should be in place. Logging of user and system actions will be crucial for any investigation.

An often-overlooked risk is **model theft** or misuse. Proprietary ML models can be reverse-engineered if not protected. Buyers may request details on model encryption or obfuscation, especially if intellectual property must be guarded.

Ethical Considerations and Bias

Key Questions: What steps has the vendor taken to identify and mitigate bias in their model? Are fairness audits conducted (e.g. checking performance across demographic groups)? Is there human-in-the-loop oversight for critical decisions?

Why It Matters: Biased AI can lead to harm: e.g., an ML model for dosing might perform poorly on underrepresented populations, endangering patient safety or reinforcing health disparities.

Actionable due diligence:

- **Bias Testing:** Consider the general algorithmic-bias processes described in [IEEE 7003-2024](#). Request documentation that the model's predictions were evaluated across relevant subgroups (age, race, sex). For instance, if the AI predicts side-effect severity, does it perform equally well on all races? Bias audits should include confusion matrices per subgroup.
- **Data Diversity:** Review the training data's demographics. If crucial subpopulations are missing, the model is likely biased.
- **Human Oversight:** For high-stakes outputs, the vendor should specify how humans review AI suggestions. Ask for the interface design or SLAs addressing human review.

The clinical-leader commentary emphasizes "algorithmic bias and fairness" as part of due diligence (www.mondaq.com). If a vendor claims a model is "industry proven" or "FDA quality," ensure those claims hold when considering bias: e.g., FDA may require evidence that a diagnostic AI was tested on a diverse clinical

trial.

Contractual and Service Considerations

Key Questions: Is the vendor stable (financially and operationally) enough to deliver long-term? What are the support and maintenance terms (patching, retraining, upgrades)? How will you ensure continuity if the vendor is acquired or exits the market?

Why It Matters: Unlike single-use lab equipment, AI is a service that often evolves. If a vendor disappears, the pharma company might be stranded with an unsupported algorithm embedded in critical processes. Contract terms must handle this:

- **Vendor Viability:** Perform company due diligence – funding, years in business, core investors. For small startups, the risk is higher; ensure escrow of code or a plan for technology transfer.
- **Service Level Agreements (SLAs):** SLAs should cover not only uptime (server availability) but also performance guarantees. For instance, an SLA might promise retention of >X% precision or cadence of model retraining.
- **Right to Audit:** Contracts should allow the pharma buyer to audit the vendor's practices periodically, including security and validation processes.
- **Data and Model Portability:** Upon contract termination, can the customer retrieve their data and model? For instance, risk-based contracts might require vendors to provide a cleaned snapshot of the AI model's learned parameters to facilitate handover.
- **Liability and Indemnity:** As noted, vendors will cap liability. Buyers should secure indemnification clauses for IP infringement and data breaches, and push for broader coverage especially if the AI decision-making could cause patient harm.

Legal advisers should review all terms carefully. The proposed contractual checklist in Table 2 summarizes questions, evidence, and acceptance criteria for negotiation.

Case Studies and Illustrative Examples

The following scenarios illustrate questions to ask during vendor due diligence.

Example: Reviewing Training and Validation Evidence

Key questions for due diligence:

- **Scrutinize Training Data:** How many training examples were used? What are their sources, annotation methods, and coverage? Do they represent the intended clinical setting?
- **Demand Independent Verification:** What independent evaluations support the exact product version and intended use? Do the studies measure recommendation concordance, clinical outcomes, or both?
- **Check Applicability:** For which populations, local formularies, and clinical guidelines was the tool tested?

Example: Generative AI for Regulatory Documents

An emerging trend is using Large Language Models (LLMs) to draft regulatory documents or medical summaries. Suppose a vendor claims their LLM can generate Approved Label changes 10× faster. A due diligence tester should:

- Request examples of actual regulatory documents drafted by the tool and independently reviewed. Check for hallucinations or unsupported claims.
- Ask what regulatory documents, if any, were used to fine-tune the model. Test the proposed model on representative documents and independently review its outputs.
- Verify control measures: e.g., does the tool include audit trails showing exactly which text was AI-generated and which was human-approved? Are draft edits tracked?

The Klick Health/Momentum Events survey collected 40 online responses between July 11 and October 8, 2025 ([survey publisher's account](#)). Refer to the qualified promotional-review findings in the adoption section above, and assess their relevance to the proposed document workflow before using them in procurement. Test the tool on the buyer's own compliance tasks and require traceable provenance for generated content.

Example: Screening for Adverse Events

Consider a vendor claiming an AI model can flag serious adverse event (SAE) reports with 95% precision, reducing case intake work by half. Due diligence would involve:

- **Data Bias:** The vendor must clarify whether their training data included balanced examples of different SAEs. If 90% of training cases are head trauma and only 1% are rare cardiovascular events, the model might overlook the rare but critical cases (high false negatives). The buyer should see confusion matrices per event type.
- **Metric Dissection:** Precision of 95% sounds good, but what about recall? If the model avoids false flags (high precision) by only catching obvious cases, it may miss many real SAEs (low recall). For patient safety, recall (sensitivity) can be more critical in pharmacovigilance. Buyers should demand both precision and recall figures ([pmc.ncbi.nlm.nih.gov](#)).
- **External Audit:** Ideally, the buyer would test the model on known datasets. For example, take 100 adjudicated SAE reports (with known outcomes) from internal archives, and evaluate model predictions. Did the model flag them correctly? Only with such in-house testing can vendor claims be validated.

As [45] emphasizes, any PV AI system “*must be validated*” under GxP ([pmc.ncbi.nlm.nih.gov](#)). Thus, simply trusting a vendor's “95%” claim without seeing the validation protocol would violate best practices.

Legal, Compliance, and Contractual Due Diligence

AI vendor claims often carry legal and compliance implications. Due diligence must therefore include a thorough review of legal risk.

Contractual Protections and Liability

AI vendors try to protect themselves; buyers need to counterbalance:

- **Liability Caps:** Many AI contracts include liability caps that are grossly insufficient for life science consequences. Clients should negotiate higher caps, especially for breaches of warranty and indemnification. For instance, if an AI diagnostic misclassification leads to patient harm, the financial repercussions can be enormous. Vendors should bear responsibility for fundamental algorithmic errors. Pharmaceutical companies often require vendors to carry professional liability insurance for such risks ([innovacc.com](https://www.innovacc.com)).
- **Indemnification:** Ensure the contract indemnifies the client from IP claims and data misuse. If an AI product inadvertently infringes a patent (e.g. by copying a patented molecule structure), the vendor should cover legal costs. Many vendors disclaim “open-source” dependencies; the contract must clarify who is liable if open-source code has a restrictive license or vulnerability.
- **Service Level Agreements (SLAs):** Standard IT SLAs (99.9% uptime, response times) are necessary but not enough. For AI, include performance SLAs tied to outcome measures. For example, “If the model’s accuracy on client-specific validation drops below X%, vendor will have 48 hours to remediate.” Without such clauses, a covert model failure could trigger business disruption with no recourse.
- **Termination and Data Retrieval:** The contract should guarantee that upon project completion or termination, the pharma company retains rights to all produced data and knowledge. If the AI involved machine-learned parameters, the contract must allow exporting model weights or retraining using the customer’s data alone (to avoid vendor “lock-in”).

Legal counsel should identify the laws that apply to the intended use and update contracts accordingly. Involve attorneys with relevant AI or healthcare experience, and require the legal basis for any claimed mandatory control.

Regulatory Compliance Clauses

Vendors should explicitly commit in contracts to follow applicable regulatory standards. For example:

- **Medical Device Regulations:** For applicable finished-device manufacturers commercially distributing medical devices, specify compliance with FDA’s [Quality Management System Regulation \(QMSR\)](https://www.fda.gov/quality/quality-management-system-regulation-qmsr), effective February 2, 2026 and incorporating ISO 13485:2016. Identify the manufacturer’s responsibilities, authorized device indications/version, and supporting quality-system and validation evidence. Do not apply this device-specific framework automatically to every pharmaceutical AI vendor. For EU devices, separately assess the applicable MDR requirements and conformity-assessment documentation.
- **Data Protection:** Contracts must specify how data is handled, including data subject rights. For GDPR-covered data, the vendor is likely a “processor,” and the contract must include standard data processing addendums. Clarify whether the vendor can use de-identified customer data for model improvement or benchmarking – or if they “shall not” use it beyond the agreed purpose ([pharmaphorum.com](https://www.pharmaphorum.com)).

One recent article emphasizes that “*end-users will want to ensure the vendor has sufficient ‘skin in the game’ while users should not take on exposure the vendor will not assume.*” ([pharmaphorum.com](https://www.pharmaphorum.com)). In practice, that means strong clauses on indemnities and warranties. Test whether the vendor is resistant to such negotiations

– resistance can be a red flag as much as a technical flaw.

Compliance With Existing Guidance

Finally, due diligence should verify whether the AI vendor already complies with or references known oversight frameworks:

- **GAMP 5 and CSV:** Identify the system's regulated activity and electronic-record scope before prescribing validation. Treat GAMP as industry guidance and distinguish it from applicable predicate-rule duties and **FDA's Part 11 recommendations**. Request a documented risk assessment, appropriate validation evidence, audit trails, and change controls for the mapped scope.
- **FDA Guidance Status:** Identify the relevant guidance, its date, scope, and draft or final status. Guidance recommendations do not themselves create a dashboard obligation. Consider performance dashboards as a proposed monitoring control, and ask counsel to identify any binding monitoring requirement separately.
- **Privacy Laws:** Ensure contract and practice align with HIPAA/HITECH (US), PIPEDA (Canada), or equivalent. For generative AI, confirm that PHI is not being fed into public LLMs inadvertently – or if it is, that Business Associate Agreements are in place. The FiercePharma article notes that many companies have outright banned ChatGPT use to avoid unsanctioned data leakage (www.fiercepharma.com).

Legal due diligence is not merely checking boxes; it is about aligning contractual risk with business risk in a complex, shifting legal environment. Engaging in-depth with legal counsel is critical here.

Security and Privacy Due Diligence

Given pharma's IP sensitivity and data protection obligations, buyers must rigorously vet vendor security.

Data Usage and Privacy

For AI vendors, data flows are often contentious:

- **Data Ownership:** The original data (e.g. patient records, chemical libraries) should unequivocally remain the property of the pharma company. Even if used to train the model, the vendor must commit that ownership does not transfer. The contract should specify that any model outputs generated belong exclusively to the client (and vice versa, the vendor owns its model's underlying code).
- **Use of Customer Data:** Buyers should restrict vendor reuse of customer data beyond the contracted service. Require documented permissions, a lawful basis, safeguards, and applicable contracts before fine-tuning. A buyer may choose a restrictive policy permitting only de-identified or synthetic training data; present that as the buyer's policy rather than a universal legal requirement. Where research PHI is involved, evaluate the relevant **HIPAA research provisions** rather than assuming consent and de-identification are always both required.
- **GenAI Prompt Data:** As highlighted in [16], generative AI vendors often claim rights over prompts or outputs. If the pharma scientist enters trade secrets or private formulas as prompts, the vendor must not repurpose them. (pharmaphorum.com). Inspect contracts for "Training Data Rights" clauses and insist

on strict limitations.

Practically, ask vendors:

- Do you store customer prompts or files after use? Distinguish unnecessary vendor retention or reuse from required regulated-record retention. FDA's [Part 11 guidance](#) preserves applicable predicate-rule retention and availability duties. Agree the authoritative archive, required source records, outputs, approvals, audit evidence, retention periods, access rights, and export arrangements. Use ephemeral vendor processing only after confirming that required records remain preserved and retrievable elsewhere; delete other data under the approved schedule.
- How do you ensure PHI is not exposed (e.g. via misconfiguration or bugs)? The vendor must have clear data deletion and retention policies.

HHS's [HIPAA Security Rule cybersecurity update](#) was issued as a proposed rule on December 27, 2024; it is not an AI-specific HIPAA rule. HHS states that the current Security Rule remains in effect during this rulemaking. For applicable covered-entity and business-associate workflows, document relevant data flows, safeguards, and risk analysis. Assess GDPR duties separately for processing within its scope; do not treat a proposal as an enforceable new obligation.

Cybersecurity Measures

High-level security questions include:

- **Standards and Audits:** Does the vendor undergo penetration testing and vulnerability assessments? Are third-party audit summaries available (SOC2 reports or ISO 27001 certificates)? Ideally, review these for any critical findings (major vulnerabilities or deficiencies).
- **Network Security:** Are all services encrypted in transit (e.g. TLS/SSL 1.2+)? Is data at rest encrypted with customer-managed keys? If the AI solution includes a web interface, does it use modern authentication (SAML, OAuth, MFA)?
- **Incident Response:** Has the vendor had any breaches? What is their incident response plan? Can the buyer get notifications within agreed timeframes if an incident occurs? The contract should define breach reporting deadlines (often 24–72 hours).
- **Supply Chain Security:** Many advanced AI tools rely on open-source libraries and third-party components. Request a Software Bill of Materials (SBOM) if possible. Ensure the vendor monitors dependencies for vulnerabilities (e.g. Spectre/Meltdown, Log4Shell) and has patching protocols.

Even if a vendor assures “we’re secure,” probing these details is essential. Sensitive IP leaks can create irreversible harm. One should treat AI vendors like any critical IT vendor – require compliance with the company’s procurement security policies (for example, hitting minimum security scorecards, as is done with cloud providers).

Performance and Integration Considerations

Beyond validating claims in isolation, a pharma buyer must examine how the AI will actually work within the organization.

Proof-of-Concept and Pilot Testing

A best practice is conducting a small-scale proof of concept (POC) or pilot using real data. This accomplishes several goals:

- **Testing Vendor Claims:** The vendor's marketing metrics can be verified on a subset of the company's own data. If a customer has historical data labeled by experts, running the AI on it can reveal performance gaps.
- **User Feedback:** Clinicians or scientists can trial the tool in a controlled environment. Their feedback will uncover usability issues early on. One pharma found that even though an AI could technically retrieve documents, scientists abandoned it when they couldn't phrase queries in the "right" way (www.pharmoutsourcing.com).
- **Integration Feasibility:** The POC often highlights practical hurdles. For instance, does the AI require data in a particular format (CSV vs. relational)? Does connecting to the corporate data warehouse require custom ETL work?

Before committing, ensure the POC has a defined "success criteria" agreed upon by both sides (e.g. "tool must correctly classify 90% of validation cases by day 10"). Document the process and results meticulously. This trial period can be part of contract negotiation (e.g. a conditional payment based on pilot success).

Change Management and User Adoption

Even a powerful AI can fail without adoption. The Human Factors of AI include:

- **Usability/User Interface (UI):** Is the tool intuitive? Pharma environments often have legacy interfaces and jargon. If an AI demands users learn new interfaces or complex commands, adoption stalls (www.pharmoutsourcing.com) (www.pharmoutsourcing.com). During due diligence, involve actual users (scientists, pharmacists, trial coordinators) and solicit their feedback on demos or mock-ups.
- **Explainability:** As noted, lack of transparency can erode trust (www.pharmoutsourcing.com). In user testing, confirm that the system's outputs are accompanied by explanations or confidence indicators that users can interpret. If an AI provides recommendations with no rationale, users are likely to reject or override them.
- **Training and Education:** Check that the vendor provides training materials tailored to pharma contexts. Who is included in training (clinical staff vs IT vs QA)? The vendor should offer robust documentation and ideally hands-on workshops. Per the Innovaccer blog, a good vendor will offer role-based training programs and support channels (innovaccer.com).

If the due diligence process surfaces apprehension (e.g. "our team fears being replaced" (www.pharmoutsourcing.com)), plan change management too. Some companies set up internal "AI champions" networks to pilot new tech and share successes, as recommended by Baringa (www.baringa.com)

(www.baringa.com). Buyers should make this part of the engagement: not only check if the AI works, but if the organization *will* work with it.

Implications and Future Directions

Evolving Standards and Oversight

As of 2026, AI in pharma remains under intense development by regulators. Buyers should anticipate more formal standards forthcoming:

- **AI Act and FDA Guidance:** The EU AI Act is enacted; high-risk classification follows the intended use and applicable product/conformity-assessment conditions, rather than pharmaceutical analytics as a blanket category. Consult the Commission's [classification guidance and application schedule](#). Its current timetable sets December 2, 2027 for high-risk rules and August 2, 2028 for AI embedded in regulated physical products. Keep a documented applicability assessment and update it before procurement or a material intended-use change; review research exclusions separately from deployed tools.
- **Industry Initiatives:** Include the general algorithmic-bias practices in [IEEE 7003-2024](#), discussed above, in the technical review where appropriate. Confirm the current designation and scope of standards cited in an RFP; assess their relevance to the intended use and record the acceptance evidence requested from the vendor.
- **Ongoing Performance Monitoring:** The view of AI as “set-and-forget” is shifting. Future due diligence will likely require systems for continuous monitoring: FDA and EMA emphasize that models should be tracked after deployment (to detect drift or data shifts) (www.pharmaceuticalonline.com). Buyers may need to partner with vendors in setting up dashboards and retraining triggers.

The Changing Vendor Landscape

Looking forward, the market for AI tools will both proliferate and consolidate. Future offerings from established technology providers may provide alternatives to AI startups. Meanwhile, industry consortiums might pool resources for data sharing, which could empower incumbents.

For the due diligence checklist, this means:

- Plan for **model and vendor evolution**: vendor claims today may become outdated quickly. For example, if the vendor relies on GPT-4 for summarization today but next year open-source domain LMs match it, future buyers will expect ability to upgrade. Contracts and diligence should allow for technological updates.
- Open Source Impact: Many companies will incorporate open-source foundation models (e.g. Meta Llama or Hugging Face models). The checklist should incorporate open-source compliance (licenses, reproducibility) as a subitem.

Ask vendors how they address prompt injection, model inversion, and other threats in the proposed deployment. When evaluating [federated learning](#), distinguish local distributed training and shared model updates from separate privacy protections: encryption, secure aggregation, and other safeguards should be

specified and tested rather than assumed. NIST cautions that federated learning alone is not a complete privacy solution.

Strategic Due Diligence

Overall, an AI vendor due diligence checklist should not be static. It must evolve with:

- **Company Strategy:** Align AI purchases with the strategic goals of the pharma enterprise. The npj framework recommends starting with strategic alignment and value case analysis (pmc.ncbi.nlm.nih.gov). If an AI solution does not clearly tie into a quantifiable problem, it may not warrant detailed scrutiny or investment.
- **Feedback and Learning:** Each AI procurement should inform the next. Successful due diligence builds organizational knowledge of “what good looks like” in AI tools. Tracking vendor performance post-deal (e.g. through KPIs) feeds back into the evaluation of future vendors.

In summary, **due diligence is an ongoing process**, interweaving technical assessment with legal review and change management. It requires significant effort, but given the stakes in pharma—patient safety, IP protection, and regulatory compliance—such rigor is warranted.

Conclusion

Pharmaceutical companies stand at a crossroads with AI: the potential benefits are vast, from faster drug discovery to more efficient trials, but so are the risks. The past few years have shown that unvetted AI can lead to wasteful spending or, worse, patient harm. A rigorous, structured due diligence process is not just prudent; it is essential to responsibly harness AI in this domain.

This report has assembled a detailed “due diligence checklist” covering technical, regulatory, security, and operational dimensions, drawing on real-world examples and expert guidance. Key themes include:

- **Treat AI as a controlled system:** rigorously validate models against intended use, just as any GMP system would be.
- **Demand transparency:** insist on clarity about data, algorithms, and testing methodology. Claims of outsize performance require independent evidence.
- **Prioritize safety and compliance:** never accept vendor claims of “compliance” without documentation, and ensure all AI usage adheres to FDA/EMA and data protection rules.
- **Manage risk contractually:** use legal tools to align incentives, including performance-based SLAs and robust indemnities.
- **Engage users early:** test tools in pilot programs and confirm alignment with workflows to ensure adoption.

In the words of one AI governance expert: *“AI vendor relationships are not set-it-and-forget-it.”*

(www.clinicalleader.com) As the regulatory landscape matures and AI technologies advance (e.g. with generative AI, edge AI, etc.), continuous vigilance will be needed. Future due diligence will likely involve new items (e.g. guardrails against synthetic data misuse, AI auditing by third parties).

Nevertheless, the core principle remains constant: **trust but verify**. With disciplined vetting of vendor claims, pharma companies can confidently integrate AI into their processes, unlocking value without sacrificing safety or compliance.

Table 1: Proposed Technical Validation Checklist

Set thresholds with domain, quality, regulatory, and security reviewers before the pilot; the criteria below are proposed procurement controls, not universal regulatory requirements.

Area	Questions	Evidence to request	Proposed acceptance criteria
Intended use and scope	What task, users, populations, and deployment are supported?	Intended-use statement, model/version inventory, applicability assessment	Named owners approve the scope and limitations before testing.
Data quality and rights	Are datasets representative, traceable, and permitted for this use?	Provenance, quality checks, permissions, split documentation	Required permissions are documented; agreed coverage and quality checks pass; leakage is excluded.
Comparative performance	How does the exact version compare with both general-purpose and specialist alternatives?	Blinded representative test protocol, held-out data, results and error analysis	Prespecified task metrics and critical-error thresholds pass; material subgroup gaps are reviewed.
Robustness and human review	What happens with missing data, unusual inputs, or unsupported outputs?	Stress tests, abstention/escalation rules, reviewer workflow	Agreed failure cases trigger a tested fallback and human review.
Security and records	Are access, data flows, audit evidence, and archives controlled?	Security assessment, audit logs, retention/export plan	Approved controls pass testing; required records remain retrievable; unresolved critical findings prevent rollout.
Change and integration	Can updates be governed and the workflow supported?	Pilot results, change-control plan, monitoring and rollback tests	Users approve the workflow; changes require agreed notification, testing, and release approval.

Table 2: Proposed Contractual Checklist

Counsel and the accountable business owners should adapt these proposed clauses to the product, jurisdiction, and allocation of responsibilities.

Topic	Questions	Evidence to request	Proposed acceptance criteria
Scope and status	What product/version and intended use are being purchased?	Contract schedule, applicable regulator decisions, responsibility matrix	The signed scope matches the tested deployment and accurately describes any authorization.
Data use and ownership	Can the vendor retain, reuse, or train on customer data?	Data-processing terms, approved purposes, ownership provisions	Authorized purposes, reuse restrictions, rights, and safeguards are explicit.
Service and performance	What happens if availability or agreed model performance deteriorates?	Metric definitions, SLAs, remediation and fallback plan	Measurable targets, escalation deadlines, remedies, and termination triggers are agreed.

Topic	Questions	Evidence to request	Proposed acceptance criteria
Updates and audits	Can the buyer assess changes and inspect relevant evidence?	Change-notification clauses, audit rights, validation deliverables	Material changes require notice and agreed assessment; relevant records are accessible.
Liability and incidents	Who handles breaches, infringement claims, or service failures?	Liability/indemnity terms, insurance evidence, incident procedure	Counsel accepts the allocation, coverage, reporting duties, and response arrangements.
Exit and retention	Can required records remain available when the contract ends?	Export format and test, archive/retention schedule, transition plan	Exit retrieval is tested; required records remain available; deletion occurs only under the approved schedule.

References

All claims and recommendations above are supported by sources from industry reports, journal articles, and expert commentaries, as cited in the text:

- [9] Korrapati, "Trust But Verify: Validating AI in Pharma's GxP World", *Pharmaceutical Online*, Oct. 2025 (discussion of AI validation principles and regulatory expectations) (www.pharmaceuticalonline.com) (www.pharmaceuticalonline.com).
- [12] Suttan, Mike, "Ultimate AI Vendor Checklist for Healthcare Leaders", *Innovaccer Blog*, Feb. 2025 (technical, domain, risk, performance due diligence aspects) (innovaccer.com).
- [14] Karofsky, E., "Why Pharma AI Projects Fail: The Human Problem Behind Technical Success", *Pharmaceutical Outsourcing*, Aug. 2025 (case series on adoption issues and UX importance) (www.pharmoutsourcing.com) (www.pharmoutsourcing.com).
- [15] Laws, L., interview with T. O'Connell, "Why generic AI fails pharma and the case for specialized AI", *Discover Pharma*, Jul. 2025 (expert discussion of domain-specific vs generic AI failures) (discover-pharma.com) (discover-pharma.com).
- [16] Pharmaphorum, "10 legal risks to consider when implementing AI in pharma supply chains" (Oct. 2025) (legal risks: data use, bias, liability, contractual safeguards) (pharmaphorum.com) (pharmaphorum.com).
- [18] O'Connor & Colgan Dunlap, "Advice on selecting your first AI-enabled vendor", *Clinical Leader*, Dec. 2025 (Q&A on terms like algorithm transparency and necessary diligence steps) (www.clinicalleader.com) (www.clinicalleader.com).
- [20] Ruggio, M., "Due Diligence In AI Vendor Selection", *Lexology/Mondaq*, Aug. 2024 (outline of due diligence focus areas: privacy, compliance, bias, liability) (www.mondaq.com).
- [27] Ross, C. & Herper, M., "AI & drug discovery: separate hype from reality", *STAT News*, May 2024 (quotes on AI hype in drug development) (www.statnews.com) (www.statnews.com).

- [30] Binkley et al., "An early pipeline framework for assessing vendor AI solutions to support ROI", *npj Digit. Med.*, Jun. 2025 (governance framework for vendor AI selection) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/46888888/)) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/46888888/)).
- [34] Reddy et al., "Unlocking the business value from AI in Pharma", *Baringa Whitepaper*, June 2025 (insights on data readiness, governance, and AI scaling) (www.baringa.com) (www.baringa.com).
- [39] HealthManagement, "Leveraging advanced methods to evaluate AI-Pharma companies", May 2019 (commentary on complexity of AI in pharma) ([healthmanagement.org](https://www.healthmanagement.org)).
- [42] Becker, Z., "AI adoption is an 'immediate priority' to most Big Pharmas, report finds", *FiercePharma*, Jul. 2025 (Define Ventures survey on AI adoption, budgets, governance structures) (www.fiercepharma.com) (www.fiercepharma.com).
- [43] Taylor, N. P., "Pharma pros skeptical of letting AI loose on regulatory compliance submissions: survey", *FiercePharma*, Nov. 2025 (survey of pharma professionals on trust in AI for regulatory files) (www.fiercepharma.com) (www.fiercepharma.com).
- [45] [47] Huysentruyt et al., "Validating Intelligent Automation Systems in Pharmacovigilance", *Drug Saf.*, Feb. 2021 (framework for AI validation in safety case processing) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/34888888/)) ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/34888888/)).
- Additional sources were consulted for definitions and context (including AI industry blogs, LinkedIn insights, and emerging regulatory proposals), but the core points above are substantiated by the citations listed.

Source links

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

1. <https://www.fiercepharma.com/marketing/ai-tech-immediate-priority-most-big-pharmas-many-plan-open-their-pockets-further-ai>
2. <https://www.statnews.com/2024/05/20/artificial-intelligence-drug-development-hype-reality-insitro/>
3. <https://www.clinicalleader.com/doc/can-t-miss-advice-on-selecting-your-first-ai-enabled-vendor-0001>
4. <https://www.pharmoutsourcing.com/Featured-Articles/621020-Why-Pharma-AI-Projects-Fail-The-Human-Problem-Behind-Technical-Success/>
5. <https://www.nature.com/articles/s41591-026-04431-5>
6. <https://intuitionlabs.ai/articles/fda-ai-drug-development-guidance>
7. <https://www.pharmaceuticalonline.com/doc/trust-but-verify-validating-ai-in-pharma-s-gxp-world-0001>
8. <https://pubmed.ncbi.nlm.nih.gov/articles/PMC7892696/>
9. <https://pharmaphorum.com/deep-dive/10-legal-risks-consider-when-implementing-ai-pharma-supply-chains>
10. <https://intuitionlabs.ai/articles/iso-27001-life-sciences-compliance-guide>
11. <https://innovaccer.com/resources/blogs/are-healthcare-organizations-evaluating-ai-vendors-correctly-before-signing>
12. <https://intuitionlabs.ai/articles/pharma-regulatory-submission-challenges>
13. <https://intuitionlabs.ai/articles/rwd-validate-trial-economic-assumptions>

14. <https://www.baringa.com/en/insights/pharma-ai/unlocking-the-business-value-from-ai-in-pharma/>
15. <https://www.ecfr.gov/current/title-21/section-211.68>
16. <https://www.businesswire.com/news/home/20251125598238/en/Survey-65-Percent-of-Pharma-Marketers-Distrust-AI-for-Creating-Regulatory-Compliance-Submissions>
17. <https://www.vitrana.com/benchmarking-ai-solutions-in-pharmacovigilance-kpis-and-best-practices/>
18. <https://www.fiercepharma.com/marketing/pharma-pros-skeptical-letting-ai-loose-regulatory-compliance-submissions-survey>
19. <https://discover-pharma.com/why-generic-ai-fails-pharma-and-how-purpose-built-ai-is-revolutionizing-drug-development/>
20. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12170878/>
21. <https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/premarket-notification-510k>
22. <https://www.fda.gov/medical-devices/premarket-submissions-selecting-and-preparing-correct-submission/de-novo-classification-request>
23. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological>
24. <https://www.fda.gov/medical-devices/digital-health-center-excellence/digital-health-software-precertification-pre-cert-pilot-program>
25. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-6>
26. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-2>
27. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
28. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/part-11-electronic-records-electronic-signatures-scope-and-application>
29. <https://www.hhs.gov/hipaa/for-professionals/security/guidance/guidance-risk-analysis/index.html>
30. https://commission.europa.eu/law/law-topic/data-protection/information-business-and-organisations/obligations_en
31. <https://www.hhs.gov/hipaa/for-professionals/special-topics/research/index.html>
32. <https://healthmanagement.org/s/leveraging-advanced-methods-to-evaluate-ai-pharma-companies>
33. <https://standards.ieee.org/ieee/7003/11357/>
34. <https://www.mondaq.com/unitedstates/new-technology/1503766/due-diligence-in-ai-vendor-selection-a-cornerstone-for-healthcare-organizations>
35. <https://www.fda.gov/medical-devices/postmarket-requirements-devices/quality-management-system-regulation-qmsr>
36. <https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/index.html>
37. <https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>
38. <https://www.nist.gov/blogs/cybersecurity-insights/privacy-attacks-federated-learning>

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.