

# AI-Enabled Clinical Research Operations

From Concept to Implementation

---

**Kush Dhody, M.D., M.S.**

President, Amarex Clinical Research, LLC (An NSF Company)

Physician-Scientist | Clinical Development Executive | Regulatory Strategy Advisor

***The question is no longer whether AI will transform clinical research. The question is whether organizations will lead the transformation or be transformed by it.***

This paper presents a strategic framework for deploying artificial intelligence across clinical research operations, from regulatory and quality review to data management, medical writing, and trial oversight. Drawing on more than two decades of clinical development leadership and direct experience implementing AI-enabled workflows in CRO operations, it examines how AI is reshaping clinical research, evaluates the emerging regulatory landscape including the FDA's January 2025 draft guidance on AI in drug development, and proposes a risk-based approach to AI adoption that balances innovation with the rigorous quality standards clinical research demands.

---

# Table of Contents

---

1. Executive Summary .....	3
2. The AI Imperative in Clinical Research .....	4
3. AI Applications Across the Clinical Trial Lifecycle .....	6
4. Practical AI Deployment in CRO Operations .....	8
5. AI-Enabled Clinical Quality and Regulatory Oversight .....	10
6. Risk-Based Credibility Assessment for AI Models .....	12
7. Data Governance and Validation in AI-Enabled Workflows .....	13
8. Organizational Change Management for AI Adoption .....	15
9. Regulatory Landscape: FDA, MHRA, and the Global Framework .....	17
10. Recommendations .....	19
11. Conclusion .....	21

---

# 1. Executive Summary

---

Clinical research is entering a structural transformation driven by artificial intelligence, one that will redefine how clinical development programs are designed, executed, and governed. Across the drug development lifecycle, AI applications are accelerating timelines, reducing costs, and enabling capabilities that were not feasible through traditional methods alone. The global AI in clinical trials market reached approximately \$2.7 billion in 2025, with projected growth to \$8.5 billion by 2030 at a compound annual growth rate of 24 to 28 percent.<sup>1</sup> By the end of this decade, AI is expected to be integrated into 60 to 70 percent of all clinical trials.

AI does not introduce new risk into clinical research. It exposes existing inefficiencies at scale.

Yet the gap between AI's promise and its responsible deployment in regulated clinical research environments remains substantial. Clinical research is not a technology sector that happens to involve patients; it is a patient safety enterprise that must evaluate every innovation against the standards of data integrity, regulatory compliance, and scientific rigor that define good clinical practice. The adoption of AI without a disciplined implementation framework introduces risks that are both measurable and consequential: algorithmic bias in patient stratification, unvalidated outputs in regulatory submissions, and erosion of the human oversight that forms the foundation of ethical clinical research.

In January 2025, the U.S. Food and Drug Administration published its first comprehensive draft guidance on the use of AI to support regulatory decision-making for drug and biological products, establishing a risk-based credibility assessment framework that signals the direction of regulatory expectations globally.<sup>3</sup> The UK Medicines and Healthcare products Regulatory Agency has set out its strategic approach to AI regulation, launching the AI Airlock regulatory sandbox and articulating five principles for proportionate oversight.<sup>9</sup> The AdvaMed white paper on AI in healthcare provides the industry's perspective on balancing innovation with patient safety.<sup>8</sup>

At Amarex Clinical Research, we have moved from concept to implementation: deploying AI capabilities across regulatory and quality review, data management, and medical writing workflows that have reduced operational turnaround times while maintaining the accuracy and compliance standards that regulatory environments demand. This experience, combined with the regulatory developments of 2024 and 2025, provides the foundation for the framework presented in this paper.

Artificial intelligence is not a tool to optimize isolated functions. It is a catalyst for re-architecting how clinical development is governed, executed, and scaled in a regulated environment. This white paper argues that the responsible adoption of AI in clinical research requires three conditions: a clear understanding of where AI creates genuine value across the trial lifecycle, a risk-based governance framework that ensures model credibility and data integrity, and an organizational change management strategy that positions AI as an augmentation of scientific expertise rather than a replacement for it.

## 2. The AI Imperative in Clinical Research

Clinical research faces structural pressures that have intensified over the past decade. Drug development costs continue to escalate, with the average cost of bringing a new drug to market now estimated at \$2.6 billion. Clinical trial timelines have lengthened: in 2024, 45 percent of pharmaceutical companies reported that their trials had grown longer over the preceding two years.<sup>2</sup> Protocol complexity has increased, with the mean number of endpoints, procedures, and eligibility criteria per study rising steadily across all phases.

### The Efficiency Gap

The consequences of this complexity are reflected in operational metrics that have resisted improvement through incremental process optimization alone:

Metric	Current State
Average Phase III trial duration	3.5 to 5 years
Clinical trial failure rate (all phases)	Approximately 90%
Protocol amendment rate	76% of Phase I–IV trials require at least one amendment
Average cost per protocol amendment (Phase III)	\$535,000 median direct cost
Patient recruitment timeline overruns	80% of trials fail to meet initial enrollment timelines
Data query resolution cycle	15 to 30 days average

Table 1. Clinical Trial Operational Benchmarks

These figures represent not merely operational inconveniences but structural barriers to the development and delivery of therapies to patients who need them. The magnitude of these challenges suggests that incremental improvements in existing workflows will not be sufficient. AI offers the potential for a different category of solution: one that addresses root causes rather than symptoms.

### Why AI, and Why Now

Three converging developments have made AI deployment in clinical research both feasible and urgent:

First, the data infrastructure required for AI has matured. Electronic data capture systems, electronic health records, real-world data repositories, and digital health technologies now generate structured datasets at a scale and granularity that AI models require. The foundation exists; what has been missing is the analytical capability to exploit it.

Second, AI technologies themselves have advanced beyond narrow applications. Natural language processing, large language models, machine learning classification, and predictive analytics have

---

demonstrated performance levels that are directly applicable to clinical research tasks. AI-driven drug discovery programs are showing clinical trial success rates of 80 to 90 percent in Phase I and 40 percent in Phase II, substantially exceeding traditional benchmarks.<sup>1</sup>

Third, the regulatory environment has evolved from cautious observation to active framework development. The FDA's January 2025 guidance, informed by review of over 500 submissions with AI components since 2016, provides the clearest signal to date that regulatory agencies are prepared to evaluate and accept AI-generated evidence when supported by appropriate credibility assessments.<sup>3,5</sup>

The strategic question facing clinical research organizations is no longer whether to adopt AI, but how to adopt it responsibly, at what pace, and with what governance structures in place.

---

## 3. AI Applications Across the Clinical Trial Lifecycle

---

AI's value is not uniform across the lifecycle; it is concentrated in decision points where data complexity exceeds human processing capacity. AI is not a single technology applied uniformly across clinical development. It encompasses a spectrum of capabilities that address distinct challenges at different stages of the trial lifecycle. Understanding this spectrum is essential for prioritizing investment and setting realistic expectations.

### Drug Discovery and Preclinical Development

While drug discovery falls outside the scope of most CRO operations, the impact of AI at this stage shapes the programs that reach clinical development. AI-powered target identification, compound screening, and toxicity prediction are accelerating the transition from laboratory to clinic. Pharmaceutical-AI partnerships surged 14-fold between 2019 and 2023, reaching \$12.8 billion in transaction value.<sup>1</sup> These programs will increasingly arrive at CROs with AI-derived data as part of their development packages, requiring CROs to evaluate and integrate AI-generated evidence.

### Clinical Trial Design and Protocol Optimization

AI models can analyze historical trial data to identify design elements statistically associated with protocol amendments, enrollment challenges, and regulatory deficiencies. Specific applications include eligibility criteria optimization using real-world data to balance scientific rigor with enrollment feasibility, endpoint selection informed by AI analysis of disease progression patterns and therapeutic response data, site selection models that predict enrollment performance based on investigator experience, patient demographics, and site infrastructure, and protocol complexity scoring that identifies design elements likely to generate amendments.

AI-powered clinical trial design has demonstrated the potential for 25 percent faster protocol approval and 30 to 50 percent fewer protocol amendments.<sup>2</sup>

### Patient Recruitment and Enrollment

Patient recruitment remains the most frequently cited operational challenge in clinical research, with nearly 30 percent of industry professionals identifying it as their primary concern.<sup>2</sup> AI addresses this through natural language processing of electronic health records to match patients against eligibility criteria, predictive models for enrollment forecasting and site-level performance, digital outreach optimization using AI-driven patient engagement platforms, and real-time enrollment monitoring with adaptive site activation strategies.

In controlled testing across 400 clinical studies, AI algorithms for clinical trial forecasting yielded an average time savings of 12 weeks compared to traditional timeline management.<sup>2</sup>

---

## **Data Management and Quality**

AI applications in clinical data management include automated data reconciliation and discrepancy detection, intelligent query generation that prioritizes clinically significant data issues over routine queries, real-time data quality monitoring with predictive identification of emerging data integrity concerns, and automated coding of adverse events, concomitant medications, and medical history using natural language processing.

## **Pharmacovigilance and Safety Monitoring**

AI is transforming safety monitoring through automated signal detection in large-scale safety databases, natural language processing of case narratives for adverse event identification and classification, predictive models for identifying patients at elevated risk of specific adverse events, and integration of real-world data sources for continuous safety surveillance beyond the controlled trial environment.

## **Regulatory Submission and Compliance**

AI capabilities are being deployed to support regulatory submission processes through automated compliance checking against regulatory requirements, generation of submission-ready documents and summaries, consistency verification across the common technical document, and predictive analytics for regulatory review timelines and potential deficiency queries.

---

## 4. Practical AI Deployment in CRO Operations

---

The transition from AI theory to operational deployment in a CRO environment requires a pragmatic assessment of where AI delivers the greatest value relative to implementation complexity and risk. Having led this transition at Amarex, I can speak directly to the operational realities.

The constraint in regulatory and quality review is not expertise; it is the mechanical burden of document handling.

### Regulatory and Quality Review

Regulatory and quality review processes represent one of the highest-value applications of AI in CRO operations. Traditional audits and quality reviews of pharmaceutical products require review of tens of thousands of documents, organized and assessed against complex international and country-specific regulations. This process has historically required four to six weeks of intensive work by teams of subject matter experts.

The deployment of AI-enabled regulatory and quality review workflows has transformed this process. At Amarex, we have implemented cloud-based AI systems that perform document inventory and completeness verification, automated organization of documents into regulated folder structures, version control tracking across large document corpora, extraction and synthesis of findings across thousands of documents into structured summaries, and compliance gap identification against regulatory requirements.

The result is a reduction in regulatory and quality review turnaround time from four to six weeks to approximately two weeks, while maintaining the accuracy and compliance standards that regulatory environments demand.

### Data Management

AI deployment in data management operations requires careful distinction between tasks where AI augments human judgment and tasks where it replaces manual processes. High-confidence automation targets include data entry verification and transcription checking, coding assistance for adverse events and concomitant medications, query generation for objective data discrepancies, and study documentation assembly and cross-referencing.

Human-in-the-loop requirements remain essential for clinical assessment of safety signals, eligibility determination for complex medical histories, data interpretation requiring therapeutic area expertise, and final quality review of AI-generated outputs.

---

Human oversight is not a temporary safeguard; it is a permanent design principle in regulated AI systems.

## Medical Writing

AI is reshaping medical writing operations, though the nature of the impact requires careful articulation. AI does not write clinical study reports or regulatory submissions autonomously. What it does is accelerate the preparatory and drafting phases of medical writing through automated extraction and organization of data from clinical databases, first-draft generation of routine document sections such as disposition tables, demographic summaries, and protocol deviation listings, consistency checking across related documents within a submission package, and literature search and summarization for background and discussion sections.

The medical writer's role evolves from assembler of information to scientific editor and strategic communicator, with AI handling the mechanical aspects of information retrieval and organization.

## 5. AI-Enabled Clinical Quality and Regulatory Oversight: A Scalable CRO Model

Regulatory and quality review processes are among the most document-intensive, time-sensitive, and consequential activities in CRO operations. They represent an ideal entry point for AI deployment, not because they are the ultimate objective, but because they provide a controlled, high-value proving ground from which AI capabilities can be extended across the broader clinical development enterprise. Our experience at Amarex implementing AI-enabled regulatory and quality review workflows provides a practical model for how CROs can move from pilot to production in a regulated environment.

### The Regulatory and Quality Review Challenge

Pharmaceutical regulatory and quality reviews typically involve tens of thousands of documents that must be accounted for, organized, reviewed, summarized, and synthesized in alignment with a wide variety of country-specific and international regulations. The traditional review process requires four to six weeks of intensive manual effort by teams of subject matter experts. The bottleneck is not expertise; it is the mechanical burden of document management, version tracking, regulatory mapping, and cross-referencing that consumes the majority of review time.

### Architecture of an AI-Enabled Review Workflow

Deploying AI in regulatory and quality review operations is not a matter of applying a single model to a single task. It requires an integrated workflow architecture with distinct AI capabilities operating in sequence:

Capability	Function	Value
Document Intelligence	Automated inventory, extraction, and classification of documents from large clinical trial archives	Eliminates manual document cataloging; ensures completeness
Regulatory Classification	Automated sorting and filing of documents into internationally regulated folder structures	Reduces classification errors; enforces consistency across audits
Version Control Automation	Automated tracking of document versions and revision histories across the archive	Prevents version conflicts; maintains regulatory-grade traceability
Intelligent Summarization	Extraction and synthesis of findings across thousands of documents into structured, audit-ready summaries	Accelerates expert review; focuses human attention on substantive findings
Compliance Gap Detection	Automated identification of missing documents, incomplete records, and regulatory non-conformities	Reduces oversight risk; prioritizes remediation efforts
Security and Access Control	Role-based access controls ensuring data access is limited to authorized personnel	Maintains data integrity and confidentiality throughout the audit

Table 2. AI-Enabled Regulatory and Quality Review Workflow Components

---

## Operational Results

The implementation of this architecture at Amarex has delivered measurable outcomes. Regulatory and quality review turnaround time has been reduced from four to six weeks to approximately two weeks, a reduction of 50 percent or more. The AI-enabled workflow achieves near-perfect accuracy in document classification, organization, and extraction. Staff experts are freed from manual, repetitive tasks to focus on higher-level scientific analysis and strategic oversight. The entire workflow operates within a secure cloud environment, with private connections and role-based access controls ensuring data confidentiality and integrity.

## Lessons from Implementation

Several lessons from this deployment are generalizable to other AI implementations in regulated environments.

First, the AI system does not replace the expert reviewer; it restructures the reviewer's work. The value proposition is not that AI performs the review, but that it removes the mechanical overhead that prevents subject matter experts from spending their time on the analysis and judgment that only they can provide.

Second, accuracy must be validated continuously, not assumed. Near-perfect accuracy in a pilot environment does not guarantee sustained performance across different sponsors, therapeutic areas, and regulatory frameworks. Ongoing monitoring and recalibration are essential.

Third, security architecture must be designed from the outset, not retrofitted. Clinical trial data flowing through AI systems must be protected against unauthorized access, modification, and loss. Role-based access controls, private network connections, and encrypted data storage are foundational requirements, not optional enhancements.

Fourth, the architecture must be modular and replicable. An AI solution designed for pharmaceutical regulatory and quality review should be architecturally capable of adaptation to medical devices, biologics, and other regulated products without fundamental redesign.

## Strategic Implications

The regulatory and quality review acceleration model demonstrates a broader principle: in highly regulated environments, the most effective AI applications are those that automate structured, high-volume, rule-governed tasks while preserving and enhancing the role of human expertise in interpretation, judgment, and decision-making. This principle applies across CRO operations, from data management to regulatory submissions.

Regulatory and quality review workflows serve as a controlled entry point. The broader implication is the transformation of regulatory, quality, and clinical oversight functions across the entire development lifecycle.

---

## 6. Risk-Based Credibility Assessment for AI Models

---

The FDA's January 2025 draft guidance, "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products," establishes the most comprehensive regulatory framework to date for evaluating AI in drug development.<sup>3</sup> Its implications extend well beyond FDA-regulated submissions to inform how any organization should govern AI deployment in clinical research.

### Scope and Applicability

The guidance applies to AI models used to produce information or data intended to support regulatory decision-making regarding safety, effectiveness, or quality for drugs. This encompasses a broad range of applications including AI to support development programs under Investigational New Drug Applications, AI in novel clinical trial designs, AI-enabled digital health technologies used in drug development, AI in pharmacovigilance, AI in pharmaceutical manufacturing, and AI in studies using real-world data to generate real-world evidence.<sup>3,6</sup>

Notably, the guidance does not address AI models used in drug discovery or those used for operational efficiencies that do not impact patient safety, drug quality, or the reliability of clinical study results.

### The Seven-Step Credibility Framework

The FDA proposes a risk-based credibility assessment framework organized in seven steps:<sup>3,7</sup>

- Step 1:** Define the question of interest that will be addressed by the AI model.
- Step 2:** Define the context of use for the AI model, specifying the role and scope of the model.
- Step 3:** Assess the risk associated with the AI model based on the context of use.
- Step 4:** Develop a credibility assessment plan proportionate to the identified risk level.
- Step 5:** Execute the credibility assessment activities, including data quality evaluation, model development documentation, and performance testing.
- Step 6:** Compile a credibility assessment report documenting findings and conclusions.
- Step 7:** Maintain and update the model as needed, with ongoing monitoring of performance.

### Implications for CRO Operations

For CROs deploying AI in clinical research operations, this framework establishes clear expectations. Every AI model used in processes that could affect regulatory decision-making must have a defined context of use. Risk assessment must be conducted before deployment, not after. Credibility activities must be proportionate to the risk level: higher-risk applications such as AI used in patient stratification or endpoint assessment require more extensive validation than lower-risk applications such as document organization.<sup>3,12</sup>

Most AI models within the scope of clinical research operations will be considered medium to high risk because they are being used in contexts that directly or indirectly affect clinical trial data, safety assessments, or regulatory submissions. Organizations must be prepared to document and defend their AI model development, validation, and maintenance processes.

## 7. Data Governance and Validation in AI-Enabled Workflows

The deployment of AI in clinical research operations creates new data governance requirements that extend beyond traditional data management frameworks. AI models are only as reliable as the data they consume, and the consequences of data quality failures in AI-enabled workflows can propagate faster and more broadly than in manual processes.

### Data Quality Requirements for AI

AI models in clinical research require data that meets standards beyond those applied to traditional analyses:

Dimension	Requirement	Consequence of Failure
Completeness	Training and input data must cover the full range of expected scenarios	Model performance degrades on underrepresented populations or edge cases
Accuracy	Source data must be verified against ground truth or validated reference standards	Model outputs inherit and potentially amplify source data errors
Consistency	Data formats, coding conventions, and terminology must be standardized	Model performance varies unpredictably across data sources
Timeliness	Training data must reflect current clinical practice and regulatory requirements	Models may generate outputs based on outdated standards
Representativeness	Data must reflect the demographic and clinical diversity of target populations	Algorithmic bias in patient-facing applications

Table 3. Data Quality Requirements for AI in Clinical Research

### Validation Framework

**Model validation** must confirm that the AI model performs as intended for its specified context of use. This includes testing on held-out datasets, evaluation across relevant subgroups, and assessment of edge case performance.

**Process validation** must confirm that the integration of AI into operational workflows maintains data integrity across the end-to-end process. This includes verification that AI outputs are correctly received, interpreted, and acted upon by downstream processes.

**Output validation** must confirm that AI-generated outputs meet the quality standards required for their intended use. This includes comparison of AI outputs against expert-generated reference standards and ongoing monitoring of output quality over time.

---

## Audit Trail Requirements

AI-enabled workflows must maintain audit trails that are at least as comprehensive as those required for manual processes. This includes documentation of model version and configuration at the time of each output generation, input data provenance and quality assessment results, model decision logic for outputs that influence regulatory decisions, human review and approval records for AI-generated outputs, and change management documentation for model updates and retraining events.

Compliance with 21 CFR Part 11 requirements for electronic records and electronic signatures remains fully applicable to AI-enabled workflows. Organizations must ensure that their AI deployment architecture supports the required controls for data integrity, access management, and audit trail maintenance.

---

## 8. Organizational Change Management for AI Adoption

---

AI adoption in clinical research is fundamentally an organizational transformation challenge, not a technology deployment exercise.

The technical challenges of AI deployment in clinical research are substantial, but the organizational challenges are often more consequential. Technology implementation without corresponding organizational transformation produces tools that are underutilized, misapplied, or resisted.

### The Change Management Imperative

Half of biopharmaceutical companies are already integrating AI and machine learning into parts of the clinical trial process or plan to do so in the near term.<sup>2</sup> Yet collaboration gaps between sponsors, CROs, and sites remain a key hurdle, and growing concern over policy and funding shifts continues to shape strategy. The gap between AI adoption intention and successful implementation is a change management problem, not a technology problem.

### A Phased Adoption Framework

Successful AI adoption in clinical research operations requires a phased approach that builds capability, confidence, and governance structures incrementally.

**Phase 1: Assessment and Foundation (Months 1 through 6).** During this phase, organizations should conduct an AI readiness assessment across technology infrastructure, data maturity, workforce capability, and governance frameworks. They should identify three to five high-value, lower-risk use cases for initial deployment. Establishing an AI governance committee with representation from clinical operations, regulatory affairs, quality assurance, data management, and information technology is essential. Developing AI-specific standard operating procedures covering model selection, validation, deployment, monitoring, and retirement is also necessary. Finally, organizations should initiate workforce training programs that address both technical AI literacy and the strategic rationale for AI adoption.

**Phase 2: Pilot Deployment (Months 6 through 12).** This phase involves deploying AI solutions in controlled pilot environments with defined success criteria and measurement frameworks. Organizations should implement human-in-the-loop workflows with clear escalation protocols for AI uncertainty, establish performance baselines and monitoring dashboards, document lessons learned and iterate on governance frameworks based on pilot experience, and begin regulatory engagement on AI deployment strategies where applicable.

**Phase 3: Scaled Implementation (Months 12 through 24).** During this phase, organizations should extend AI deployment to additional use cases based on pilot results and organizational readiness. They should integrate AI workflows into standard operational processes with appropriate quality controls, develop cross-functional AI competency centers that serve as internal centers of excellence, establish continuous improvement processes for AI model performance and governance, and build partnerships with technology providers and regulatory consultants to maintain alignment with evolving best

practices.

## Workforce Transformation

AI adoption does not eliminate the need for clinical research professionals; it transforms their roles. The evolution follows a consistent pattern across functions:

Function	Current Role	AI-Augmented Role
Clinical Data Manager	Manual data review, query generation, coding	AI oversight, complex data interpretation, quality strategy
Medical Writer	Document assembly, data extraction, formatting	Scientific editing, strategic communication, AI output review
Clinical Research Associate	Source data verification, documentation review	Risk-based oversight, AI-flagged issue investigation, site partnership
Regulatory Affairs Specialist	Manual compliance checking, document preparation	Strategic regulatory planning, AI-assisted submission review
Quality Assurance Auditor	Document-by-document review, manual findings compilation	AI-directed risk-based auditing, systems-level quality analysis

Table 4. Workforce Role Evolution with AI Adoption

The critical success factor is framing AI as augmentation of scientific expertise rather than displacement of professionals. Organizations that position AI as a threat to existing roles will face resistance that undermines adoption. Organizations that position AI as a tool that elevates professional work from mechanical tasks to strategic contributions will build the engagement necessary for successful transformation.

## 9. Regulatory Landscape: FDA, MHRA, and the Global Framework

The regulatory environment for AI in clinical research is evolving rapidly across multiple jurisdictions. Organizations deploying AI must maintain awareness of this evolving landscape and build governance frameworks flexible enough to accommodate regulatory developments as they occur.

### FDA: From Discussion to Framework

The FDA’s engagement with AI in drug development has followed a deliberate trajectory:

Date	Milestone
December 2022	Expert workshop convened by Duke Margolis Institute for Health Policy on behalf of CDER/FDA
May 2023	Discussion paper published on AI use in drug development; received over 800 external comments
August 2024	Public workshop on guiding principles for responsible AI use in drug and biological product development
January 2025	Draft guidance: “Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products”
January 2025	Separate draft guidance on AI-enabled medical device software functions

Table 5. FDA AI Regulatory Development Timeline

The January 2025 draft guidance represents the culmination of years of engagement. Its risk-based credibility assessment framework was informed by the FDA’s experience with over 500 submissions containing AI components from 2016 to 2023.<sup>3,5</sup> The guidance applies to AI models across the drug product lifecycle: nonclinical, clinical, post-marketing, and manufacturing phases.

The CDER AI Council, established as a decisional body, coordinates AI-related activities across the center, promoting consistency in both internal AI capabilities and external regulatory policy. This institutional infrastructure signals that AI oversight is not a temporary initiative but a permanent function within the FDA’s regulatory architecture.

### MHRA: Principles-Based Regulation

The UK Medicines and Healthcare products Regulatory Agency has adopted a principles-based approach organized around five strategic principles: safety, security, and robustness; appropriate transparency and explainability; fairness; accountability and governance; and contestability and redress.<sup>9,10</sup>

The MHRA’s AI Airlock regulatory sandbox, launched in 2024, represents an innovative approach to addressing the unique regulatory challenges posed by AI as a Medical Device. By creating a controlled environment where real-world AI products are evaluated alongside regulators, the NHS, and Approved

---

Bodies, the MHRA is building regulatory knowledge through practical experience rather than theoretical frameworks alone.<sup>9</sup>

The MHRA's Data Strategy 2024 through 2027 explicitly addresses the role of AI and machine learning in regulatory operations, including evaluation of natural language processing for pharmacovigilance systems and establishment of cross-functional task forces to explore the operational potential of generative AI and large language models.<sup>10</sup>

## International Convergence

The International Medical Device Regulators Forum released in January 2025 a final document identifying ten guiding principles for Good Machine Learning Practice, building on the guiding principles released in October 2021 by the FDA, Health Canada, and the MHRA.<sup>11</sup> This trilateral cooperation, now expanded to an international framework, demonstrates a trajectory toward regulatory convergence on fundamental AI principles.

For clinical research organizations operating globally, the convergence of regulatory frameworks around risk-based, principles-driven AI governance is a positive development. It suggests that a single, well-designed internal governance framework can be adapted to meet requirements across major regulatory jurisdictions, reducing the compliance burden of multi-jurisdictional AI deployment.

## Industry Perspective

The AdvaMed white paper on AI in healthcare provides a comprehensive industry perspective, highlighting that AI-enabled medical devices have been regulated by the FDA for over 25 years and emphasizing that the vast majority of current AI/ML products function as diagnostic tools that assist clinicians in decision-making rather than making independent decisions.<sup>8</sup> The paper advocates for risk-proportionate regulation that avoids creating barriers to innovation while maintaining patient safety standards.

---

## 10. Recommendations

---

Drawing on two decades of clinical development experience and direct experience deploying AI in CRO operations, I offer the following recommendations for organizations navigating the transition to AI-enabled clinical research operations:

### **Start with Structured, High-Volume Regulatory and Quality Processes**

The highest-value, lowest-risk entry point for AI in CRO operations is document-intensive, rule-governed regulatory and quality review processes where AI augments rather than replaces expert judgment. Regulatory and quality review workflows, data reconciliation, and document management provide measurable returns while building organizational confidence and governance maturity.

### **Adopt the FDA's Risk-Based Framework as an Internal Standard**

Even for organizations not immediately subject to the FDA's January 2025 draft guidance, the seven-step credibility assessment framework provides a rigorous and defensible governance structure. Organizations that adopt this framework proactively will be prepared for regulatory requirements as they formalize across jurisdictions.

### **Invest in Data Infrastructure Before AI Models**

AI model performance is constrained by data quality. Organizations should prioritize data standardization, governance, and infrastructure investments before or concurrent with AI model deployment. A sophisticated AI model operating on poorly governed data will produce sophisticated errors.

### **Build Human-in-the-Loop Workflows as a Design Principle**

In clinical research, AI outputs that affect patient safety, data integrity, or regulatory decisions must be reviewed by qualified human experts. This is not a limitation of current technology that will be resolved by future advances; it is a permanent design principle for regulated environments. Build workflows with human oversight designed in from the outset, not retrofitted as an afterthought.

### **Treat Change Management as Equal to Technology Implementation**

Allocate the same strategic attention, budget, and leadership engagement to organizational change management as to technology procurement and deployment. The organizations that succeed with AI will be those that transform their culture, not just their technology stack.

### **Engage Regulators Early and Often**

The FDA, MHRA, and other regulatory agencies have explicitly encouraged early engagement on AI deployment strategies. Take them at their word. Early dialogue reduces regulatory risk, shapes realistic expectations, and builds the institutional relationships that facilitate efficient review of AI-enabled submissions.

---

## **Govern for the Long Term**

AI governance is not a one-time compliance exercise. Establish governance structures, including an AI governance committee, standard operating procedures, validation frameworks, and monitoring systems, that are designed for continuous operation and evolution as both technology and regulatory expectations advance.

---

# 11. Conclusion

---

AI is not arriving in clinical research; it has arrived. The question confronting sponsors, CROs, and regulatory agencies is not whether AI will be deployed, but whether its deployment will be governed by deliberate strategy or reactive improvisation.

Across more than two decades leading global clinical development programs and CRO operations, I have seen the industry navigate transformations before: the digitization of clinical data, the globalization of clinical trials, the adoption of risk-based monitoring. Each transformation followed a similar pattern. Early movers who invested in understanding the change, built appropriate governance frameworks, and engaged regulators proactively captured disproportionate value. Organizations that waited for mandates or adopted technology without strategic intent found themselves in perpetual catch-up.

AI is the next transformation, and it is the most consequential. Its impact will be felt across every function of clinical development, from the first protocol draft to the final regulatory submission. Our experience deploying AI-enabled regulatory and quality review workflows at Amarex demonstrates what is achievable today: measurable operational improvements delivered within a framework of regulatory rigor and data security. The FDA's January 2025 guidance provides the regulatory architecture for responsible deployment.<sup>3</sup> The technology, the regulatory framework, and the operational proof points are in place.

What remains is leadership: the willingness to invest, to govern, and to transform. The organizations that lead will not be those with the largest AI budgets or the most advanced algorithms. They will be the organizations that deploy AI with the same discipline, rigor, and commitment to patient safety that has always defined excellent clinical research.

AI does not replace judgment. It amplifies it. And in an industry where judgment determines whether patients receive therapies that improve their lives, that amplification is not merely valuable. It is imperative.

The organizations that understand this will not just adopt AI. They will define how it is used.

---

## **Kush Dhody, M.D., M.S.**

President, Amarex Clinical Research, LLC (An NSF Company)

Physician-Scientist | Clinical Development Executive | Regulatory Strategy Advisor

[kushdhody.com](http://kushdhody.com)

---

# References

---

1. Clinical Leader. "Global AI in Clinical Trials Market Trends, Current Partnerships." February 2025. <https://www.clinicaleader.com/doc/global-ai-in-clinical-trials-market-trends-current-partnerships-0001>
2. PPD/Thermo Fisher Scientific. "The AI Revolution in Clinical Trials." Market Report, 2024. [https://www.ppd.com/wp-content/uploads/2025/11/Market-Report\\_The-AI-revolution-in-clinical-trials.pdf](https://www.ppd.com/wp-content/uploads/2025/11/Market-Report_The-AI-revolution-in-clinical-trials.pdf)
3. U.S. Food and Drug Administration. "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products." Draft Guidance, January 2025. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological>
4. U.S. Food and Drug Administration. "Using Artificial Intelligence & Machine Learning in the Development of Drug & Biological Products." Discussion Paper, May 2023 (Revised February 2025). <https://www.fda.gov/media/167973/download>
5. U.S. Food and Drug Administration. "FDA Proposes Framework to Advance Credibility of AI Models Used in Drug and Biological Product Submissions." Press Release, January 6, 2025. <https://www.fda.gov/news-events/press-announcements/fda-proposes-framework-advance-credibility-ai-models-used-drug-and-biological-product-submissions>
6. Goodwin Law. "FDA Publishes Draft Guidance on Use of Artificial Intelligence in the Drug Product Life Cycle." January 2025. <https://www.goodwinlaw.com/en/insights/publications/2025/01/alerts-lifesciences-aiml-fda-publishes-its-first-draft-guidance>
7. DLA Piper. "Key Takeaways from FDA's Draft Guidance on Use of AI in Drug and Biological Life Cycle." January 2025. <https://www.dlapiper.com/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai>
8. AdvaMed. "The Role of Artificial Intelligence (AI) in Healthcare." White Paper, 2024. <https://www.advamed.org/wp-content/uploads/2024/11/AdvaMed-AI-White-Paper-Final.pdf>
9. MHRA. "MHRA's AI Regulatory Strategy Ensures Patient Safety and Industry Innovation into 2030." April 2024. <https://www.gov.uk/government/news/mhras-ai-regulatory-strategy-ensures-patient-safety-and-industry-innovation-into-2030>
10. MHRA. "Impact of AI on the Regulation of Medical Products." Strategic Report, 2024. [https://assets.publishing.service.gov.uk/media/662fce1e9e82181baa98a988/MHRA\\_Impact-of-AI-on-the-regulation-of-medical-products.pdf](https://assets.publishing.service.gov.uk/media/662fce1e9e82181baa98a988/MHRA_Impact-of-AI-on-the-regulation-of-medical-products.pdf)
11. U.S. Food and Drug Administration. "Good Machine Learning Practice for Medical Device Development: Guiding Principles." IMDRF Final Document, January 2025. <https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles>
12. Foley & Lardner LLP. "AI Drug Development: FDA Releases Draft Guidance." January 2025. <https://www.foley.com/insights/publications/2025/01/ai-drug-development-fda-releases-draft-guidance/>

---

**Disclaimer:** *The views expressed in this white paper are those of the author and do not necessarily represent the official position of Amarex Clinical Research, LLC (An NSF Company) or any regulatory authority. This document is intended for informational and educational purposes and should not be construed as regulatory, legal, or compliance advice.*

© 2025 Kush Dhody. All rights reserved.

*This white paper may be shared and cited with appropriate attribution to the author. For permissions beyond citation or excerpting, please contact the author.*