

As AI Accelerates in Life Sciences, Governance Must Keep Up



ARTICLE 07.08.26 GUI COZZI

What life sciences organizations deploying AI need to know about governance, risk, and regulatory accountability – before an incident forces the conversation.

Artificial intelligence is no longer a future consideration for life sciences. It's already embedded in how your organization operates – accelerating drug discovery, streamlining clinical documentation, powering regulatory submissions, and optimizing manufacturing quality control. The efficiency gains are real. So are the risks.

The challenge isn't whether to deploy AI. The challenge is whether your governance structure can keep pace with the deployment. For most life sciences organizations, the honest answer is: not yet.

The Regulatory Environment Didn't Wait for You

Life sciences operate under some of the most demanding compliance frameworks in existence – FDA 21 CFR Part 11, GxP, HIPAA, and increasingly, the EU AI Act's high-risk AI provisions that directly implicate clinical decision support, diagnostic tools, and automated manufacturing processes. These frameworks were designed for a world where humans made decisions and systems recorded them. AI flips that model. Systems now make – or heavily influence – decisions, and humans ratify them.

Regulators are already asking questions your organization may not be ready to answer:

- Who is accountable when an AI-assisted clinical decision leads to an adverse event?
- How do you demonstrate model integrity and data lineage during an FDA audit?
- What controls prevent your AI agents from taking autonomous actions on validated systems?
- Has your AI vendor been evaluated for privacy controls and model training data practices?

These are not hypothetical. FDA's emerging guidance on AI/ML-based Software as a Medical Device (SaMD) explicitly anticipates continuous learning models that change behavior post-deployment. The EU AI Act classifies several life sciences AI applications as high-risk, with mandatory conformity assessments. And plaintiffs' attorneys are building AI liability theories faster than most compliance teams are building AI policies.

The Incident That Should Be on Every CIO's Radar

In April 2026, an AI coding agent destroyed an automotive SaaS company's entire production database – including 90 days of backups – in nine seconds. No cyberattack. No malicious insider. Just an AI agent with unconstrained permissions, no guardrails, and no one who had formally defined what it was and wasn't allowed to do.

Now apply that scenario to a life sciences context. Your AI agent has access to your LIMS. Or your validated manufacturing execution system. Or your clinical trial data repository. What are its permission boundaries? Who defined them? Where is that documented?

What Governance-First Looks Like in Life Sciences

AI governance in life sciences isn't a checkbox exercise. It's a structured program that connects executive accountability, documented controls, and technical assurance — and maps cleanly to the regulatory and audit expectations already governing your environment.

Ownership and Decision Rights. Someone in your organization needs to own AI risk. Not informally. Formally – with defined authority, an escalation path to the Board, and a Security Steering Committee that includes clinical, legal, IT, and compliance representation. AI incidents are leadership failures when no one owns the outcome.

Acceptable Use That Actually Covers AI. Most Acceptable Use Policies were written before generative AI existed. They don't address which AI tools employees may use with patient data, how IP and proprietary research data are protected when entered into commercial LLMs, or what happens when a staff member uses a shadow AI tool to summarize a clinical report. That policy needs to be updated – specifically, not generically.

Vendor Risk for AI Vendors. Your standard vendor risk assessment wasn't designed to evaluate whether a vendor's model is trained on customer inputs, whether a productivity AI retains PHI, or whether an AI-powered CRO tool meets your data residency requirements. Third-party AI vendors need an assessment framework built for how AI actually works – covering model training data, privacy controls, security posture, and contractual data use restrictions.

AI Risk Assessment Against a Recognized Framework. The NIST AI Risk Management Framework provides a structured approach to evaluating AI risk across governance, data, model, and operational domains. For life sciences organizations, this maps well to existing quality and risk management culture — and produces the kind of documented, auditable output that regulators and auditors can evaluate.

Adversarial Testing of Your LLMs. If your organization has deployed or integrated large language models – for clinical summarization, regulatory writing assistance, or internal knowledge retrieval — those models need to be tested the way your other systems are tested: adversarially. Prompt injection, data leakage, model abuse, and evasion techniques aligned to MITRE ATLAS give you visibility into how those models behave under attack before an adversary discovers it first.

Infrastructure Controls for Agentic AI. This is the category that's moving fastest and has the least governance attention. AI agents – workflows that autonomously take actions in your environment – require a specific set of controls: least-privilege permissions, confirmation guardrails for consequential actions, audit logging, backup architecture, and environment segmentation. In a validated GxP environment, agentic AI without these controls is a compliance and patient safety risk.

The Questions Leadership Should Be Asking Now

Before your next Board meeting, executive team discussion, or audit cycle, consider whether your organization can answer these confidently:

1. Who is our AI Risk Owner, and what is their authority?
2. Does our Acceptable Use Policy explicitly govern AI tools and data handling?
3. Have we assessed our AI vendors for model training data practices and privacy controls?
4. Have we performed adversarial testing on any LLMs in our environment?
5. Do our AI agents operate within defined permission boundaries with documented guardrails?
6. Can we produce evidence of AI governance structure for an FDA auditor or a Board inquiry?

If any of these surfaces uncertainty, that's the right place to start.

What a Structured Engagement Looks Like

Dean Dorton's AI Security Governance program is a 12–16 week engagement that takes organizations from governance gap to defensible program. It covers all six components – AI Governance, Acceptable Use Policy, Vendor Risk Management, AI Risk Assessment (NIST AI-RMF), LLM Penetration Testing (MITRE ATLAS), and Infrastructure & Controls Assessment – in a phased approach tailored to your regulatory environment and operational readiness.

For life sciences clients, we integrate directly with your existing compliance and quality management infrastructure, align deliverables to audit and regulatory expectations, and ensure that governance outputs are Board-ready, not just IT-ready.

The Cost of Waiting

The organizations that will navigate the next wave of AI-related regulatory scrutiny, litigation, and incidents with confidence are building governance structures now – before they're compelled to. In life sciences, where the stakes include patient safety, validated systems, and federal regulatory relationships, the cost of a reactive posture is simply too high.

The question isn't whether your organization needs AI governance. It's whether you build it on your terms or someone else's.

Gui Cozzi is Director of Cybersecurity Risk & Compliance at Dean Dorton, a Top 100 professional services firm. He leads the firm's cybersecurity practice, including AI Security Governance, Fractional CISO services, and compliance advisory across healthcare, life sciences, financial services, and higher education.

Ready to put governance around your AI program? [Schedule a consultation.](#)