SIERRA SOLUTIONS

AI GOVERNANCE IN A REGULATED LIFE SCIENCES ENVIRONMENT

WHITE PAPER

From Policy to Practice:
Building an Enterprise AI Governance Framework

sierrasolutionsgroup.com

Big Picture Overview

A commercial-stage life sciences company engaged Sierra Solutions to build a comprehensive AI governance framework from the ground up — establishing the policies, procedures, training programs, and oversight structures needed to adopt artificial intelligence responsibly across its enterprise. As AI tools proliferated rapidly across the organization's commercial and operational functions, leadership recognized that ungoverned adoption posed meaningful risks to data privacy, regulatory compliance, and IP integrity. Sierra Solutions partnered with internal IT security and executive leadership to transform that risk into a structured, auditable, and scalable foundation for responsible AI use.

The Situation

A growing life sciences organization providing technology-enabled solutions to pharmaceutical and healthcare organizations found itself at a critical inflection point. As generative AI tools like ChatGPT, Microsoft Copilot, and specialized analytics platforms became widely accessible, employees across commercial operations, data, and clinical teams began adopting them organically — often without IT oversight, formal approval, or awareness of the data privacy implications.

Leadership faced a defining question: how do you harness the productivity benefits of AI without exposing sensitive customer data, proprietary systems, or regulated information to unacceptable risk? The answer required more than a policy — it required an enterprise-wide framework that could govern AI from the ground up.

The Challenge

The organization's challenges were multilayered and time-sensitive:

- Shadow AI proliferation — employees using unauthorized, free-tier AI tools with company credentials and sensitive data
- No formal intake or approval process for evaluating AI tools against security, privacy, and compliance requirements
- Absence of employee awareness — most staff lacked a baseline understanding of AI risk, hallucination, data leakage, or IP exposure
- Vendor AI enablement — existing third-party platforms were quietly adding AI capabilities, creating unreviewed risk in already-contracted tools
- Regulatory exposure — as a company operating in the life sciences space, non-compliant AI use carried potential HIPAA, IP, and data sovereignty implications

Our Approach

Sierra Solutions embedded with the organization's IT Security and executive leadership team to design and operationalize a full-spectrum AI governance program. The engagement was structured across four interconnected workstreams:

- AI Governance Policy Development
- AI Governance Committee & Intake Process
- Enterprise AI Training Program
- Vendor AI Risk Re-Assessment

Framework for Operational Excellence

AI Governance Policy Development

Authored a formal AI Acceptable Use Policy governing all generative AI, LLM, and machine learning tool use across employees, contractors, and consultants — aligned to NIST AI Risk Management Framework standards.

The policy was structured around tiered risk categories — distinguishing fully contracted enterprise AI platforms, vendor-embedded AI features within existing approved systems, and unauthorized public tools — each with clear usage boundaries and escalation paths. It directly addresses the highest-exposure scenarios: entering client or patient-adjacent data into public AI tools, using AI-generated output without human validation in compliance-sensitive work, and defines audit-logging expectations to keep AI usage defensible under regulatory review.

The policy was shaped through structured input from Legal, IT Security, Compliance, and frontline commercial teams — ensuring it reflected real operational workflows rather than a generic governance template, improving adoption and enforceability from day one.

AI Governance Committee & Intake Process

Established a formal four-step AI tool approval process requiring business case submission, security and privacy risk assessment, vendor use-case validation, and executive sponsor sign-off before any tool could be used with company data.

Sierra Solutions designed a cross-functional AI Governance Committee — spanning IT Security, Compliance, Legal, and business unit leadership — to serve as the standing decision-making body for all AI-related requests:

- Step 1 — Business Case Submission: Requestors document use case, data types, and expected value via a standardized intake form.
- Step 2 — Security & Privacy Risk Assessment: IT Security evaluates data residency, encryption, and access control criteria.
- Step 3 — Vendor Use-Case Validation: Vendor AI terms and sub-processor relationships are reviewed to confirm company data won't train third-party models.
- Step 4 — Executive Sponsor Sign-Off: A named executive accepts accountability for the tool's ongoing compliant use.

This structure gives the organization a single source of truth for every AI tool in use — closing the visibility gap that previously allowed shadow AI adoption to go unchecked.

The Facts

Workforce Coverage

100%

ALL EMPLOYEES, CONTRACTORS & CONSULTANTS TRAINED AND GOVERNED UNDER THE NEW FRAMEWORK

AI Approval Process

4 STEP

FORMAL INTAKE, ASSESSMENT, VENDOR VALIDATION, AND EXECUTIVE SIGN-OFF BEFORE ANY AI TOOL USE

Open Source AI Tools

0

POLICY PROHIBITS ALL OPEN-SOURCE AI; ONLY FULLY CONTRACTED ENTERPRISE VERSIONS PERMITTED

Framework for Operational Excellence

Enterprise AI Training Program

Designed and delivered an organization-wide AI Introduction Training covering AI fundamentals, the 5 Cs of AI Risk (Controls, Compliance, Confidentiality, Continuity, Culture), real-world risk case studies, and acceptable use procedures.

The training opens with foundational AI literacy — how large language models generate output, and why hallucination and data leakage occur — replacing blind trust with informed judgment. Its core is the 5 Cs of AI Risk, a practical model employees can apply in the moment:

- Controls — Is this tool approved, and what boundaries apply?
- Compliance — Does this use align with regulatory obligations?
- Confidentiality — Could this input expose proprietary or client data?
- Continuity — Is there a human fallback if output is wrong or unavailable?
- Culture — Does this reflect the organization's values around accountability?

Real-world case studies of public AI data leakage make the risks concrete, with scenario-based exercises tailored to each function. Delivery combined live sessions for leadership and high-risk roles with asynchronous modules — driving 100% completion across employees, contractors, and consultants.

Vendor AI Risk Re-Assessment

Implemented an ongoing process to re-evaluate all existing third-party vendors that had introduced AI capabilities into their platforms — ensuring previously approved tools remained compliant under the new governance framework.

A key risk Sierra Solutions identified: existing, already-vetted vendors were rolling out new AI features independently — often with no notice or re-contracting — meaning a platform approved years earlier under traditional security criteria could suddenly process company data in entirely new ways.

To close this gap, Sierra Solutions built a standing re-assessment process, including a full inventory of vendor contracts flagged for AI presence, a recurring review cycle to catch newly introduced capabilities, a standardized AI risk questionnaire, and a re-approval workflow routing flagged vendors back through the Governance Committee.

This ensures governance doesn't stop at point-in-time approval — giving the organization a repeatable way to catch vendor-introduced AI risk before it becomes unmanaged.

HYPE VS. REALITY

✗ Hype: "Adopt AI now, govern it later."	✗ Hype: "A policy document is enough."	✗ Hype: "Our vendors already have this covered."	✗ Hype: "This is an IT problem."
✓ Reality: Ungoverned adoption is why most AI pilots never scale.	✓ Reality: Policy without process is just a PDF nobody follows.	✓ Reality: Vendors add AI features overnight — oversight has to be ongoing.	✓ Reality: Every employee is a potential AI risk — and a potential safeguard.

Strategic Impact & Business Value Delivery

The establishment of a proactive, enterprise-wide AI governance framework has transformed artificial intelligence from an unmanaged operational risk into a governed strategic capability. By building the policy infrastructure, oversight mechanisms, and workforce readiness programs needed for responsible AI adoption, the organization is now positioned ahead of the regulatory curve — with the governance architecture in place to scale AI confidently as the technology and compliance landscape continues to evolve.

Business Impact

Eliminated Shadow AI Risk

A formal acceptable use policy and four-step approval process closed the gap on unauthorized AI tool use across the organization. Employees and contractors now operate under clear boundaries — with written standards governing what tools can be used, with what data, and under what conditions — significantly reducing the organization's exposure to data leakage, IP loss, and compliance violations.

Established Regulatory Readiness

With the EU AI Act's high-risk AI provisions taking full effect in August 2026 and the FDA's first AI guidance for drug and biological products finalized in 2025, the organization's governance framework directly anticipates emerging regulatory expectations — reducing the compliance burden of future mandates before they become enforcement priorities.

Built a Culture of AI Accountability

The enterprise-wide AI Introduction Training program established a shared vocabulary and risk awareness baseline across all teams. The 5 Cs framework — Controls, Compliance, Confidentiality, Continuity, and Culture — gave every employee a practical model for evaluating AI use in their own role, replacing ad hoc experimentation with informed, accountable adoption.

Created Scalable Governance Infrastructure

The AI Governance Committee, vendor re-assessment process, and formal intake workflow are designed to scale as the organization's AI footprint grows. Governance keeps pace with adoption — ensuring that as new tools emerge and existing vendors add AI capabilities, each is evaluated consistently before reaching production environments.

The Result

Through its strategic partnership with Sierra Solutions, this commercial-stage life sciences organization has built the governance infrastructure to adopt AI responsibly at enterprise scale. By authoring a comprehensive acceptable use policy, standing up a formal AI Governance Committee, delivering organization-wide AI literacy training, and establishing ongoing vendor risk oversight, the engagement delivered more than compliance — it delivered confidence. Leadership now has the oversight structures and workforce readiness needed to pursue AI-driven productivity gains without the operational and regulatory risks that have derailed similar initiatives at peer organizations. The governance framework established through this engagement is not a point-in-time deliverable; it is a living system designed to evolve alongside the technology, the regulatory environment, and the organization's own strategic ambitions.