

Establishing a Generative AI Governance Framework and Scoping Use Cases for a Biopharma Company

INDUSTRY:



Pharma

PRODUCTS AND SERVICES:



Therapeutics

BUSINESS TRANSFORMATION STORY & BUSINESS CHALLENGE

The client, a biopharmaceutical company, was preparing to roll out Microsoft Copilot across the enterprise. Employees were eager to adopt generative AI, but the organization needed a practical path to move from AI readiness uncertainty to a governed, scalable adoption model.

At the start of the engagement, the company had no policy defining acceptable use. It also lacked a consistent way to assess risk or determine where AI could safely touch clinical, regulatory, or patient-related data. Without clear governance and guardrails, the company faced two competing risks: an uncontrolled rollout that could expose regulated information or an overly cautious approach that stalled adoption and limited the value of its Copilot investment.

The company also needed a structured way to identify where Copilot could create the most operational value. That meant evaluating use cases across the business, prioritizing opportunities based on value and risk, and sequencing adoption in a way that supported reasonable growth.

Clarkston Consulting partnered with the client to establish the governance need for enterprise-wide AI adoption, equip employees to use Copilot responsibly, and develop a value-driven adoption roadmap that aligned business priorities with a phased rollout.

PRIMARY OBJECTIVES

- Establish a generative AI governance framework defining acceptable use, risk tiers, human oversight requirements, and incident response for a regulated life sciences environment
- Define clear boundaries for AI use around the company's regulated and patient-related data
- Enable the workforce to use Microsoft Copilot confidently and safely through role-relevant training grounded in the company's own governance framework, including how to judge whether a task is appropriate for AI, how to validate output, and how to escalate when something goes wrong
- Identify and prioritize the highest-value, lowest-effort use cases to turn the company's Copilot investment into measurable operational efficiency
- Build an enduring operating model the company could sustain independently as AI capability matured





RESOLUTION

- Drafted a four-document governance framework operating as one connected system: an Acceptable Use of Generative AI Tools policy, an AI Use-Case Risk Matrix and Review Guide, a Human-in-the-Loop Validation Guide, and an AI Incident Escalation Playbook
- Designed a tiered risk model that classified AI use from generally acceptable to restricted-by-default and prohibited, with restricted categories covering patient data, clinical trial data, regulatory submissions, quality and GxP records, pharmacovigilance information, and autonomous agents
- Established the human-in-the-loop (HITL) principle anchoring the entire framework: AI assists, humans own the output
- Developed and delivered a comprehensive Microsoft Copilot training program for approximately 70 employees, covering core capabilities, prompting technique, advanced application-specific prompts across Word, Excel, PowerPoint, Outlook, and Teams, agent creation, and the full governance framework
- Built practical decision tools translating policy into daily practice, including a go/no-go decision flow for evaluating whether a task is appropriate for Copilot, a visual risk ladder mapping each tier to a clear action, and incident-response guidance
- Designed and facilitated an AI Use Case and Business Value Workshop with function-specific templates, a structured intake process, and follow-ups with each function to refine their submissions
- Developed a weighted use-case prioritization model that scored each opportunity on business value and implementation effort, accounting for build complexity and data readiness, and translated the submissions into a decision-ready backlog mapped by quick wins, strategic bets, easy adds, and deprioritized items
- Analyzed the completed backlog for overlap and duplication and found that the requests resolved into a small set of reusable capabilities rather than discrete projects, including structured data extraction, meeting-content to deliverable workflows, clinical review and synthesis, analysis and scientific writing, financial close and commentary, and everyday self-serve prompting

KEY BENEFITS

- Established a complete generative AI governance foundation, moving the company from no formal AI policy to a defined and documented operating model for AI use in a regulated environment
- Enabled approximately 70 employees to use Microsoft Copilot with a clear understanding of both its capabilities and its boundaries
- Protected the company's most sensitive assets by defining explicit guardrails around regulated data categories before enterprise AI adoption began
- Engaged five business functions (IT, Clinical Operations, Finance, Translational Research, and Project Management) in structured use-case discovery, and delivered a scored backlog of 29 use cases sequenced from quick wins to strategic bets, with recommended starting pilots identified
- Identified a smaller set of reusable AI capabilities behind the 29 requests, helping the company shift from one-off automation ideas to scalable capability building across functions

KPIs

- 70 employees trained on Copilot capabilities, governance expectations, and safe use decision-making
- 5 business functions engaged in structured use-case discovery and prioritization
- 29 AI use cases assessed, scored, and prioritized at close of engagement
- 10 quick wins and 17 strategic-bet use cases identified for development